

THE BOUNDARY OF AN AFFINE INVARIANT SUBMANIFOLD

MARYAM MIRZAKHANI AND ALEX WRIGHT

1. INTRODUCTION

Motivation and main result. Let $\mathcal{H}(\kappa)$ be a stratum of the moduli space of Abelian differentials, i.e. the space of pairs (X, ω) where X is a Riemann surface of genus g and ω is a holomorphic 1-form on X whose zeros have multiplicities $\kappa_1, \dots, \kappa_s$ and $\sum_{i=1}^s \kappa_i = 2g - 2$. The form ω defines a flat metric on X with conical singularities at the zeros of ω . Such (X, ω) are called *translation surfaces*, and can be also described using polygons up to cut and paste.

The $GL(2, \mathbb{R})$ orbit of any translation surface equidistributes in some affine invariant submanifold, which governs most dynamical and geometric problems on the surface. Besides the obvious ones such as the entire moduli space, affine invariant submanifolds parameterize families of translation surfaces with exceptional flat geometry and dynamics, give rise to amazing families of algebraic curves, and witness unlikely intersections in the moduli space of principally polarized Abelian varieties. Each affine invariant submanifold is a rare gem, and their classification is now a central problem in dynamics on moduli spaces with rich connections to many other areas of mathematics.

For the connection to flat geometry and dynamics see, for example, [EM01, Vee89, LW15, Wri15a, AD16, CE15], and for the connection to Abelian varieties see, for example, [McM03, McM06a, Möl06b, Möl06a, Fil16, Fil, KM, MZ].

There are only countably many affine invariant submanifolds [EMM15, Wri14], and we hope for strong classification results based on inductive arguments using degenerations of translation surface.

Such arguments should take place in a partial compactification $\overline{\mathcal{H}}$ of the stratum (moduli space of translation surfaces) compatible with degenerations obtained by expressing translation surfaces in terms of polygons and letting edge lengths go to zero. In the introduction we will

restrict to the subspace $\overline{\mathcal{H}}_{\text{conn}} \subset \overline{\mathcal{H}}$ of connected (single component) translation surfaces.

Let $\partial\mathcal{M}$ denote the boundary of \mathcal{M} in $\overline{\mathcal{H}}$. By Eskin-Mirzakhani-Mohammadi [EM, EMM15], $\partial\mathcal{M} \cap \overline{\mathcal{H}}_{\text{conn}}$ is a finite union of affine invariant submanifolds \mathcal{M}_i . The main result in this paper is to compute their tangent spaces.

Suppose \mathcal{M}_i is contained in the stratum $\mathcal{H}_i \subset \overline{\mathcal{H}}$. At points in \mathcal{H} near \mathcal{H}_i , we show that the tangent space of the boundary stratum \mathcal{H}_i is naturally identified with a subspace of the tangent space to \mathcal{H} .

Theorem 1.1 (Main Theorem, special case). *Let \mathcal{M}_i be one of the finitely many affine invariant submanifolds whose union is $\partial\mathcal{M} \cap \overline{\mathcal{H}}_{\text{conn}}$. The tangent space of \mathcal{M}_i is given by the intersection of the tangent space to \mathcal{M} with the tangent space to the boundary stratum.*

This formula is valid for a sequence of points of \mathcal{M} converging to the boundary stratum \mathcal{H}_i . See Section 2 for more precise statements, which allow multicomponent degenerations. The multicomponent case is more complicated, in part because closed $GL(2, \mathbb{R})$ invariant sets of multicomponent surfaces are not always finite unions of affine invariant submanifolds, and some of our results in the multicomponent case are conditional on a version of [EM, EMM15] for multicomponent surfaces (see Conjecture 2.10).

Two important invariants of an affine invariant submanifold are rank and field of definition [Wri15a, Wri14]. The rank is a measure of size related to the flexibility of disjoint sets of cylinders on surfaces in \mathcal{M} , and the field of definition is related to the arithmetic complexity of the linear equations locally defining \mathcal{M} . (This is not the same thing as the field of definition of an algebraic variety, which is related to the polynomial equations globally defining the variety.) As an example we mention that closed $GL(2, \mathbb{R})$ orbits have rank 1, and have field of definition \mathbb{Q} if and only if they contain a square-tiled surface.

Corollary 1.2 (Consequences of Main Theorem, special case). *Each \mathcal{M}_i has smaller dimension than \mathcal{M} and rank at most that of \mathcal{M} . The field of definition of \mathcal{M}_i is equal to that of \mathcal{M} .*

Our partial compactification $\overline{\mathcal{H}}$ will be described in more detail in Section 2. It can be obtained as the bundle of finite area Abelian differentials over the Deligne-Mumford compactification, modulo zero area subsurfaces, with zeros and possibly other points marked. For genus g surfaces, this compactification will include surfaces of all genera at most g , and hence the boundary of an affine invariant submanifold will typically consist of many pieces of different genera.

Applications. If \mathcal{M} is an affine invariant submanifold of genus g translation surfaces, then the rank of \mathcal{M} is at most g . If its rank is equal to g , it is called full rank.

Theorem 1.3. *Any affine invariant submanifold of full rank is equal to a connected component of a stratum, or the hyperelliptic locus in a connected component of a stratum.*

The proof of Theorem 1.3 appears in [MW] and uses Theorem 1.1 and other results of this paper. Theorem 1.3 has already been used in [HW].

Paul Apisa has used Theorem 1.1 and other results of this paper in his recent proof that all orbit closures of rank greater than 1 in hyperelliptic connected components of strata arise from covering constructions [Api].

Idea of proof and additional results. The collection of results we develop to prove Theorem 1.1 may be as interesting as Theorem 1.1 itself.

Our strategy starts with the Cylinder Deformation Theorem [Wri15a], which appears as Theorem 4.1 below. Every translation surface has infinitely many cylinders [Mas86]. Shearing a collection of parallel cylinders while leaving the rest of the surface unchanged gives a path of translation surfaces called a cylinder twist. The Cylinder Deformation Theorem states that cylinder twists given by equally shearing the full set of cylinders in some direction on a translation surface (X, ω) remain in any affine invariant submanifold containing (X, ω) .

In an effort to understand affine invariant submanifolds geometrically, one could consider, for each (X, ω) , the span $E(X, \omega)$ of derivatives of such cylinder twists. This $E(X, \omega)$ is a subspace of the tangent space to the stratum at (X, ω) . By the Cylinder Deformation Theorem, the dimension of $E(X, \omega)$ is a lower bound for the dimension of any affine invariant submanifold containing (X, ω) . However this $E(X, \omega)$ is very behaved: In Remark 6.1 we observe that it is nowhere continuous.

Nonetheless, we will show that the tangent space to the boundary \mathcal{M}_i is generated by cylinder twists that can still be performed on nearby surfaces in the interior \mathcal{M} . This implies that, once the tangent space to the boundary stratum \mathcal{H}_i is identified with a subspace of the tangent space to the stratum \mathcal{H} , the tangent space to the boundary affine invariant submanifold \mathcal{M}_i is contained in the tangent space to \mathcal{M} .

To do this, we will have to define a larger class of twists than those generating $E(X, \omega)$, by using knowledge of \mathcal{M} . Our larger class of

twists need not remain in all affine invariant submanifolds containing (X, ω) , only \mathcal{M} .

Recall that two cylinders on a surface $(X, \omega) \in \mathcal{M}$ with core curves α, β are said to be \mathcal{M} -parallel if there is a linear equation $\int_\alpha \omega = c \int_\beta \omega$ for some $c \in \mathbb{R}$ that holds locally on \mathcal{M} . The Cylinder Deformation Theorem states that cylinder twists given by equally shearing an equivalence class of \mathcal{M} -parallel cylinders remain in \mathcal{M} .

Suppose that a surface $(X, \omega) \in \mathcal{M}$ has exactly two cylinders in some direction, but their ratio is not in the set of ratios of \mathcal{M} -parallel cylinders. The Cylinder Deformation Theorem then guarantees that the twist in each one of these individual cylinders remains in \mathcal{M} . Such twists will be the prototypical examples of a class of twists that we call recognizable, a notion that depends \mathcal{M} . The key properties of recognizable twists, defined in Section 6, are

- (1) at every point of \mathcal{M} , the recognizable twists remain in \mathcal{M} ,
- (2) away from a very small set (a finite union of proper affine invariant submanifolds) the recognizable twists span the tangent space to \mathcal{M} ,
- (3) recognizable twists at the boundary of \mathcal{M} are also recognizable at nearby surfaces in the interior.

To obtain these properties, we use a version of the following result.

Theorem 1.4 (Cylinder Finiteness Theorem, special case). *For each affine invariant submanifold \mathcal{M} , the set of ratios of circumferences of \mathcal{M} -parallel cylinders on surfaces in \mathcal{M} is finite.*

See Section 5 for a more powerful version, which additionally describes moduli and cylinder deformations.

The idea of the proof of the Cylinder Finiteness Theorem is to use the $GL(2, \mathbb{R})$ action and the results of Minsky-Weiss on recurrence of horocycle flow [MW02] to reduce the problem to studying cylinders of bounded circumference in a compact subset of the moduli space of unit area surfaces. This simple plan runs up again serious technical difficulties, which we resolve using new results of independent interest on cylinder deformations. For example, we provide a partial converse to the Cylinder Deformation Theorem, for which we additionally make use of results of Smillie-Weiss [SW04].

Theorem 1.5. *All cylinder deformations supported on an equivalence class of \mathcal{M} -parallel cylinders are, up to purely relative cohomology classes, equal to a scalar multiple of the standard cylinder deformation.*

The standard cylinder deformation is the deformation that twists all cylinders equally; this is precisely the cylinder deformation which is guaranteed to stay in \mathcal{M} by the Cylinder Deformation Theorem.

After we introduce the appropriate notation, Theorem 1.5 is restated as Theorem 4.7 in more precise language.

Corollary 1.6. *If the tangent space to \mathcal{M} contains no purely relative cohomology classes, then the moduli of \mathcal{M} -parallel cylinders are rational multiples of each other.*

The corollary can be seen as generalization of one part of the Veech dichotomy [Vee89]. More generally, one can typically obtain more complicated rational linear relations among moduli of \mathcal{M} -parallel cylinders when the tangent space contains some purely relative cohomology classes.

In the course of our study, we must prove basic results on the partial compactification $\overline{\mathcal{H}}$ and how strata of small genus translation surfaces appear in the boundary of strata of large genus translation surfaces, for example Proposition 9.3. Previous results primarily focused on special cases when the difference of the genus of the boundary surface and the interior surface is at most 1 [Boi15, KZ03, EMZ03].

Potential applications. The results of this paper seem to greatly restrict the structure of affine invariant submanifolds. For example, they contributed significantly to the discovery, announced in early 2014, of a rank two affine invariant submanifold of $\mathcal{H}(6)$ with field of definition $\mathbb{Q}[\sqrt{5}]$. (At the time, the results of the present paper were conjectured to be true by the authors.) We hope the results of this paper will be similarly be useful in the future for the classification of affine invariant submanifolds, either with certain properties or in certain strata. It is conceivable that more new affine invariant submanifolds could be discovered in the process.

The authors are currently pursuing generalizations of Theorem 1.3. It would be interesting to classify non-arithmetic rank 2 orbit closures in $\mathcal{H}(6)$, as we have implied might be possible above. We hope that recent efforts to classify higher rank affine invariant submanifolds in low genus might be extended using the new tools in this paper [NW14, ANW16, AN].

It is worth noting that Theorem 1.1 is useful even without any classification of the affine invariant submanifolds that might appear in the boundary. This theme appears already in the work of Apisa [Api], and will also feature in forthcoming work of the authors.

The boundary structure of affine invariant submanifolds is of interest in the dynamics of Teichmuller geodesic flow [**AMY13, Gad**]. The simple boundary structure of closed $GL(2, \mathbb{R})$ orbits (i.e., there are finitely many cusps) played a key role in the Veech dichotomy [**Vee89**], and in the recent work of Avila and Delecroix on weak mixing in Veech surfaces [**AD16**].

Our methods, in particular Theorem 6.9, may clarify the extent to which the boundary of an affine invariant submanifold determines the affine invariant submanifold.

Notes and references. For an introduction to orbit closures of translation surfaces, see, for example, [**Wri16, Wri15b, MT02, Zor06**]. For recent results on the classification of affine invariant submanifolds, see [**NW14, ANW16, AN**]. For the earlier classification results of McMullen in genus 2, see [**McM07, McM06b, McM05**]. See also [**Cal04**] for a different presentation of affine invariant submanifolds in genus 2. For some extensions of McMullen’s techniques beyond genus 2, see [**HLM12, HLM09, Ngu11**].

A number of works, some ongoing, define compactifications of strata that remember much more information than ours, for example [**EKZ14, Gen, Che, BCG⁺**] and work in progress of John Smillie. Rather than deliberately forgetting the information of vanishing subsurfaces, they individually rescale the metric on each such subsurface to obtain meromorphic differentials. Due to the meromorphic differentials the boundary of an affine invariant submanifold in such a compactification will not be an affine invariant submanifold.

Using disjoint methods from ours, Filip has shown that affine invariant submanifolds are quasi-projective varieties [**Fil16, Fil**]. It would be interesting to relate his work to the boundary theory of affine invariant submanifolds. On the one hand, one should be able to deduce some results on the boundary from Filip’s work. Filip shows that \mathcal{M} can be defined by endomorphism and torsion conditions, and it would be interesting to explicitly understand the relationship between these conditions and those defining the boundary of \mathcal{M} . This might allow a description of the boundary in algebro-geometric terms. On the other hand, Möller suggested several years ago that, given sufficiently good understand of the boundary of \mathcal{M} , one might hope to find a proof of algebraicity using Chow’s Theorem [**Mum95, Corollary 4.6**].

Although we describe the boundary of an affine invariant submanifold of \mathcal{M} , we do not describe what \mathcal{M} looks like in a neighborhood of the boundary. This interesting problem has not even been addressed

for strata: so far there is no explicit description of a neighborhood of a boundary stratum, except in special cases [KZ03, EMZ03].

Organization. In Section 2 we define our partial compactification and state our main theorem on boundaries. The proofs of basic facts about the partial compactification are deferred until the end of the paper in Section 9. In Section 3 we give examples of convergent sequences of translation surfaces. In Section 4 we recall the Cylinder Deformation Theorem, and prove the partial converse. In Section 5 we state and prove the Cylinder Finiteness Theorem. This is the technical core of the paper. Next, in Sections 6 and 7 we define and study recognizable twists, and use them to study the boundary of affine invariant submanifolds. Most of the paper is written for the case of connected (single component) degenerations, and the general case is addressed in Section 8.

Acknowledgements. We thank Matt Bainbridge, Alex Eskin, Simion Filip, Howard Masur, Curt McMullen, Kasra Rafi, Rick Schoen, Scott Wolpert, and Anton Zorich for helpful conversations. We are grateful to Curt McMullen for sharing the example in Section 3.3 with us. We thank the referee for their helpful comments and suggestions.

The work of MM was partially supported by NSF and Simons grants. The work of AW was partially supported by a Clay Research Fellowship.

2. MULTICOMPONENT SURFACES AND STATEMENT OF MAIN RESULT

In this section we introduce a natural class of finite area limits of translation surfaces and state our main results. We also review some definitions and some of our main tools.

2.1. The “what you see is what you get” partial compactification.

We begin with the definitions.

Definition 2.1. A *multicomponent translation surface* is a collection (X, ω, Σ) . Here X is a compact Riemann surface with at most finitely many connected components, ω is an Abelian differential that is nonzero on every connected component of X , and $\Sigma \subset X$ is a finite set of marked points. We require that Σ contain all the zeros of ω .

We will also consider labeled multicomponent surfaces, which are multicomponent surfaces where the components are labelled. (The labeling is to allow us refer to the different components unambiguously.)

Every abelian differential on a connected Riemann surface defines a multicomponent translation surface as above, by marking the zeros.

Marked points where ω does not vanish will be considered to be zeros of ω of order zero.

A partition is a set $\kappa = \{k_1, \dots, k_s\}$, where the k_i are non-negative integers (some of which may be zero). Typically we will assume $\sum k_i = 2g - 2$ for some integer g , and will say that κ is a partition of $2g - 2$.

Given a finite collection of partitions $\kappa_1, \dots, \kappa_p$, the stratum $\mathcal{H}(\kappa_1) \times \dots \times \mathcal{H}(\kappa_p)$ of labeled multicomponent surfaces will be the collection of all labeled multicomponent translation surfaces with p components, each one with zeros and marked points given by κ_i . The unlabeled stratum is a quotient of the labeled stratum, in which the labeling is forgotten. Frequently we will pass from the unlabeled situation to the labeled situation without comment by taking the pre-image under this forgetful map.

Definition 2.2. Let \mathcal{H}' and \mathcal{H} be two strata of multicomponent translation surfaces. Say that $(X_n, \omega_n, \Sigma_n) \in \mathcal{H}'$ converges to $(X, \omega, \Sigma) \in \mathcal{H}$ if there are decreasing neighborhoods $U_n \subset X$ of Σ with $\cap U_n = \Sigma$ such that the following holds. There are maps $g_n : X \setminus U_n \rightarrow X_n$ that are diffeomorphisms onto their range, such that

- (1) $g_n^*(\omega_n)$ converges to ω in the compact open topology on $X \setminus \Sigma$,
- (2) the injectivity radius of points not in the image of g_n goes to zero uniformly in n .

Injectivity radius at a point is defined as the sup of ε such that the ball of radius ε in the flat metric centered at that point is embedded and does not contain any marked points.

Given a stratum of connected translation surfaces, we will work in the partial compactification obtained by the union of that stratum together with all limit points of sequences of surfaces in the stratum. It follows directly from Definition 2.2 that this space admits a continuous $GL(2, \mathbb{R})$ action.

Remark 2.3. This partial compactification can also be obtained as follows. This description is not required for the main results in this paper.

- (1) Given the stratum of connected translation surfaces, mark all zeros of the Abelian differentials to obtain a subset of the Hodge bundle over the moduli space of Riemann surfaces with marked points.
- (2) Take the subset of the Hodge bundle over the Deligne-Mumford compactification consisting of all finite area Abelian differentials in the closure. (This is endowed with the subspace topology.)
- (3) Consider two Abelian differentials to be equivalent if they are equal after removing all zero area components (including nodes

that lie on zero area components), and filling in any punctures thus created with marked points. (So if a zero area component was joined at a node to another component, that node becomes a marked point.) Take the quotient by this equivalence relation. (This quotient space is endowed with the quotient topology, which is Hausdorff.)

Definition 2.2 is the same as convergence in this space. This notion of convergence was also used by McMullen in [McM13].

So as not to distract from our main results, we defer the proof of the final claim in Remark 2.3 and the next three results to Section 9.

Proposition 2.4. *Suppose that $(X_n, \omega_n, \Sigma_n) \in \mathcal{H}'$ converges to $(X, \omega, \Sigma) \in \mathcal{H}$. Then there are “collapse maps” f_n from X_n to X (with some additional identifications between points of Σ) that map marked points to marked points and have the following property: for any $\varepsilon > 0$ and n large enough, the f_n are the inverse to g_n on the subset of $(X_n, \omega_n, \Sigma_n)$ with injectivity radius at least ε .*

Related collapse maps also appear in the theory of the Deligne-Mumford compactification, see for example [Bai07].

A sequence $(X_n, \omega_n, \Sigma_n)$ may converge to a limit with more than one component. In this case one can think that some of the marked points of each component of the limit are naturally identified to some of the marked points in the other components. We do not remember these identifications in this paper, however they are required in the previous propositions, since of course there cannot be continuous maps f_n from the connected X_n to X when X is disconnected.

In fact we will not use the f_n much, only the “space of vanishing cycles”

$$V_n = \ker(f_n : H_1(X_n, \Sigma_n, \mathbb{C}) \rightarrow H_1(X, \Sigma, \mathbb{C})).$$

Note that the relative homology group $H_1(X, \Sigma, \mathbb{C})$ is unchanged if some points of Σ are identified to some other points of Σ , thus the map on relative homology exists even if the identifications among points of Σ are forgotten. Although the f_n are not canonically defined on the subset of with small injectivity radius, the induced map on relative homology does not depend on the choice of map f_n satisfying the conditions of Proposition 2.4.

Proposition 2.5. *V_n is well defined and eventually constant in n : there is a neighborhood of (X, ω, Σ) in \mathcal{H}' that contains $(X_n, \omega_n, \Sigma_n)$ for n large enough on which the space of vanishing cycles is well defined and invariant under parallel transport.*

In this paper we will not need the results of Rafi on the comparison between flat and hyperbolic metrics [Raf07], however those familiar with these results can anticipate the source of vanishing cycles. The $(X_n, \omega_n, \Sigma_n)$ have a thick part with large injectivity radius, and a thin part with small injectivity radius. These two parts are separated by a collection of “expanding annuli”. The space of vanishing cycles is the relative homology of the thin part.

Since the space of vanishing cycles V_n is eventually constant, we will often omit the n and write V instead of V_n . It will be implicit that only large n are considered. We will write $\text{Ann}(V)$ for the space of cohomology classes that annihilate the vanishing cycles, using the usual pairing between relative cohomology and relative homology.

Proposition 2.6. *In the situation above, $\text{Ann}(V)$ is naturally identified with the tangent space to the boundary stratum \mathcal{H} . Furthermore, there is a neighborhood of 0 in $\text{Ann}(V)$ such that if ξ_n, ξ are in this neighborhood and $\xi_n \rightarrow \xi$, then $(X_n, \omega_n, \Sigma_n) + \xi_n$ converges to $(X, \omega, \Sigma) + \xi \in \mathcal{H}$.*

The above result is quite intuitive: the ways of deforming the limit surface correspond to the ways to deforming the pre-limit surfaces that do not change the parts of the surface that are degenerating. Recall that strata are locally modeled by relative cohomology, hence our notation $+\xi$ to indicate the surface whose corresponding relative cohomology class is modified by adding ξ . To prove the proposition, one must in particular show that $(X_n, \omega_n, \Sigma_n) + \xi$ is well defined for all sufficiently small ξ independent of n .

2.2. Statement of results. Recall that an affine invariant submanifold is a properly immersed connected submanifold locally defined in period coordinates by real linear equations with zero constant term. We also use this definition for strata of multicomponent translation surfaces. A foundational result on affine invariant submanifolds is the following, from [EM, EMM15].

Theorem 2.7 (Eskin-Mirzakhani-Mohammadi). *Any closed $GL(2, \mathbb{R})$ invariant subset of a stratum of connected translation surfaces is a finite union of affine invariant submanifolds.*

The tangent bundle $T(\mathcal{M})$ of an affine invariant submanifold is naturally a subbundle of H_{rel}^1 , the vector bundle with fibers $H^1(X, \Sigma, \mathbb{C})$. Let H^1 denote the vector bundle with fibers $H^1(X, \mathbb{C})$, and let $p : H_{\text{rel}}^1 \rightarrow H^1$ denote the natural map.

As the next example shows, affine invariant submanifolds may have self-crossings.

Example 2.8. Let $\mathcal{M} \subset \mathcal{H}(2,0,0)$ be the locus where at least one of the two marked points is at a Weierstrass point. (The marked points are not allowed to coincide, or to be equal to the zero.) Then \mathcal{M} self-crosses along the locus where both marked points are Weierstrass point. In one branch, one marked point stays a Weierstrass point, and in the other branch the other point stays a Weierstrass point. (McMullen has also told us a simple way to build affine invariant submanifolds with self-crossings without using marked points.)

We suppress some minor issues arising from self-crossings, whose solutions can be found in [LNW, Section 2.1]. However, we keep in mind that boundary affine invariant submanifolds may have points of self-crossing, and possibly also distinct components that cross each other.

Theorem 2.9 (Main Theorem). *Let \mathcal{M} be an affine invariant submanifold.*

(1) *The subset of the boundary of \mathcal{M} consisting of connected surfaces consists of a finite union of affine invariant submanifolds. If $(X_n, \omega_n) \in \mathcal{M}$ converge to a connected translation surface (X, ω) in the boundary, then (X, ω) is contained in a component \mathcal{M}' of the boundary of \mathcal{M} whose tangent space can be computed, for infinitely many n , as*

$$T(\mathcal{M}') = T_{(X_n, \omega_n)}(\mathcal{M}) \cap \text{Ann}(V).$$

Here $V = V_n$ is the space of vanishing cycles, which is eventually constant.

(2) *If (X', ω') is a point in the boundary of \mathcal{M} , and π is a projection onto one of its connected components, then $\pi(X', \omega')$ is contained in an affine invariant submanifold \mathcal{M}' whose tangent space can be computed, for infinitely many n , as*

$$T(\mathcal{M}') = \pi_*(T_{(X_n, \omega_n)}(\mathcal{M}) \cap \text{Ann}(V)).$$

The restriction to infinitely many n , rather than all n , reflects the possibility that several components of the boundary might cross at (X, ω) , necessitating passing to a subsequence. Because the boundary is a *finite* union of affine invariant submanifolds, Theorem 2.9(1) implies that the sequence (X_n, ω_n) can always be decomposed into a union of finitely many disjoint subsequences, for each of which $T_{(X_n, \omega_n)}(\mathcal{M}) \cap \text{Ann}(V)$ is eventually equal to the tangent space of a branch of the self-crossing.

The first part of Theorem 2.9 is a more precise restatement of Theorem 1.1.

Given affine invariant submanifolds $\mathcal{M}_i, i = 1, 2$ of strata \mathcal{H}_i of possibly multicomponent surfaces, one can obtain a product affine invariant submanifold $\mathcal{M}_1 \times \mathcal{M}_2$ in the stratum of labelled multicomponent surfaces $\mathcal{H}_1 \times \mathcal{H}_2$. If \mathcal{M} does not arise in this way, we say that it is irreducible. (If \mathcal{M} is in a stratum of unlabeled surfaces, we say it is irreducible if the connected components of its pre-image in the stratum of labelled surfaces is irreducible.) Examples of irreducible affine invariant submanifolds can be obtained by considering loci of multicomponent surfaces all of whose components cover a common single component translation surface.

Conjecture 2.10 (Multicomponent EMM). Let \mathcal{M} be a closed $GL(2, \mathbb{R})$ invariant subset of a stratum of multicomponent surfaces, and assume \mathcal{M} consists entirely of surfaces all of whose components have equal area. Then \mathcal{M} is equal to the set of surfaces all of whose components have equal area in a finite union of affine invariant submanifolds. Furthermore:

- (1) There are only countably many closed $GL(2, \mathbb{R})$ invariant subsets of the locus where all components have equal area.
- (2) For any irreducible affine invariant submanifold, locally the absolute periods of any component determine those of all other components. In particular, the ratio of areas of the components is constant.
- (3) The projection of any affine invariant submanifold to any subset of the set of components is equal to the complement of a finite union of affine invariant submanifolds in a larger affine invariant submanifold.
- (4) The bundle H^1 over any affine invariant submanifold is semisimple.

We believe that if any version of the main statement could be shown for multicomponent surfaces, all four following statements would also follow immediately. The four statements can be thought of as likely corollaries of the main statement, whose proofs are deferred until the main statement has been verified.

It is plausible that the proofs in [EM, EMM15] could be adapted to the multicomponent case with only very few changes.

Remark 2.11. It is certain that at least some minor changes would be required to the proofs, because the “diagonal” action (g_t, g_t) of geodesic flow on the product of two strata is part of a \mathbb{R}^2 action, namely (g_s, g_t) . Thus, even when restricting to the set where both components have equal area, the flow (g_t, g_t) is not hyperbolic.

Theorem 2.12. *The first statement of Theorem 2.9 holds without the connectivity assumption (i.e., multicomponent limits are allowed) if multicomponent EMM is true.*

Remark 2.13. For multicomponent surfaces, not all closed $GL(2, \mathbb{R})$ invariant sets are finite unions of affine invariant submanifolds, and there are uncountably many affine invariant submanifolds. Both issues are created by the ability to scale the different components separately. These issues are mild, but add a few complications to the proof of Theorem 2.12.

The body of the paper is written in the case of connected limits. Section 8 addresses the multicomponent case, and contains the proofs of the second statement of Theorem 2.9 and the proof of Theorem 2.12. From now until Section 8, all statements, unless we explicitly say otherwise, are for the connected case.

2.3. A few corollaries. We now turn to the proof of Corollary 1.2, which is restated in more detail below.

Recall that the rank of an affine invariant submanifold is defined as $\frac{1}{2} \dim p(T(\mathcal{M}))$ [Wri15a]. (Recall $p : H_{\text{rel}}^1 \rightarrow H^1$ was defined before Theorem 2.9.) The field of definition of \mathcal{M} is the smallest subfield of \mathbb{R} such that \mathcal{M} can be defined by linear equations with coefficients in this field [Wri14].

Corollary 2.14. *Suppose that \mathcal{M} is an affine invariant submanifold, and let \mathcal{M}' be a component of the boundary of \mathcal{M} .*

- (1) *The dimension of \mathcal{M}' is less than that of \mathcal{M} .*
- (2) *The field of definition of \mathcal{M}' is equal to that of \mathcal{M} .*
- (3) *The rank of \mathcal{M}' is at most that of \mathcal{M} .*
- (4) *If $\ker(p) \cap T(\mathcal{M}) = \{0\}$, then the rank of \mathcal{M}' is strictly less than that of \mathcal{M} .*

The corollary is also true without projections assuming multicomponent EMM. All parts except the second are true for projections to connected components in the multicomponent case (without assuming multicomponent EMM). When projections are considered, one can only show that the field of definition of \mathcal{M}' is contained in that of \mathcal{M} , and indeed there are examples that show it may be strictly smaller.

Proof. All of these follow from Theorem 2.9.

- (1) This follows from $T(\mathcal{M}') = T_{(X_n, \omega_n)}(\mathcal{M}) \cap \text{Ann}(V)$. The tangent space $T(\mathcal{M})$ cannot be entirely contained in $\text{Ann}(V)$ because the $GL(2, \mathbb{R})$ directions are not contained in $\text{Ann}(V)$.

- (2) The containment $\mathbf{k}(\mathcal{M}') \subset \mathbf{k}(\mathcal{M})$ follows from the same formula, since $\text{Ann}(V)$ is defined over \mathbb{Q} . The other containment follows, for example, from [Wri14, Theorem 5.1]. Indeed, $T_{(X_n, \omega_n)}(\mathcal{M}) \cap \text{Ann}(V)$ is defined over $\mathbf{k}(\mathcal{M}')$, and so in particular $T_{(X_n, \omega_n)}$ contains vectors defined over $\mathbf{k}(\mathcal{M}')$. By [Wri14, Theorem 5.1], $\mathbf{k}(\mathcal{M}) \subset \mathbf{k}(\mathcal{M}')$.
- (3) This follows from the same formula and the following commutative square.

$$\begin{array}{ccc}
 H^1(X_n, \Sigma_n) & \xleftarrow{\quad} & H^1(X, \Sigma) \\
 \downarrow p & & \downarrow p \\
 H^1(X_n) & \xleftarrow{\quad} & H^1(X)
 \end{array}$$

- (4) This follows from part 1 and the fact that the dimension is at least twice the rank with equality if $\ker(p) \cap T(\mathcal{M}) = \{0\}$.

□

2.4. Degenerating cylinders. The following summarizes some facts that will be used throughout the paper. Say two annuli are isotopic if their core curves are isotopic.

Lemma 2.15. *Suppose $(X_n, \omega_n, \Sigma_n)$ converges to (X, ω, Σ) , and let $g_n : X \setminus U_n \rightarrow X_n$ be the maps given in Definition 2.2. The following statements hold.*

- (1) *For every cylinder C on (X, ω, Σ) and for n large enough there is a corresponding cylinder C_n on $(X_n, \omega_n, \Sigma_n)$ isotopic to $g_n(C)$. The height, modulus, and circumference of C_n converge to those of C .*
- (2) *If C_n are cylinders on $(X_n, \omega_n, \Sigma_n)$ of circumference uniformly bounded above and below and height uniformly bounded below, then, possibly after passing to a subsequence, there is a cylinder C on (X, ω, Σ) such the circumference and height and direction of C_n converges to those of C , and such that C_n is isotopic to $g_n(C)$.*
- (3) *In the situation of (2), for any $\varepsilon > 0$, for n large enough depending on ε , if C'_n is a cylinder on $(X_n, \omega_n, \Sigma_n)$ that is parallel to C_n but not isotopic to $g_n(C')$ for a cylinder C' on (X, ω, Σ) parallel to C , then either C'_n has modulus less than ε or circumference less than ε (or both).*

In fact, all of the isotopies in this result can be assumed to be small, i.e., not to move points very much.

The first claim is standard, and boils down to the fact that the image of a cylinder under a map which almost preserves the flat metric will contain a cylinder of similar size: a related claim is used in [MT02, Proposition 4.5], and it is also helpful to compare to the proof of [McM13, Theorem A.2]. The second claim is similar, and the third follows from the second.

3. EXAMPLES OF CONVERGENCE TO LOWER GENUS SURFACES

In this section we give four examples of convergence according to Definition 2.2. The first shows that the definition is easy to verify for the most commonly used degenerations. The second is actually a special case of the first, and is the simplest situation in which the degeneration cannot be obtained via local surgeries, contrary to the commonly studied situation for “splitting a zero” and “bubbling off a cylinder” [KZ03]. The last two examples show that extremely complicated behavior is possible along converging sequences.

3.1. Degenerating cylinders. Let \mathcal{C} be a collection of parallel cylinders on a multicomponent translation surface (X, ω, Σ) . Assume the union $\cup\mathcal{C}$ of the cylinders in \mathcal{C} do not cover all of (X, ω, Σ) . Let $(X_t, \omega_t, \Sigma_t)$, $t \in (0, 1)$ be the surfaces obtained by linearly scaling the height of the cylinders in \mathcal{C} by a factor of t without shearing the cylinders. Define $(X_0, \omega_0, \Sigma_0)$ as follows. First cut out the cylinders in \mathcal{C} , to obtain a surface with boundary. Mark every point p in the boundary if it is in Σ or if there is a vertical line in $\cup\mathcal{C}$ from p to a point in Σ . These finitely many marked points cut the boundary into finitely many edges. Identify two of the edges when they can be joined by vertical lines in $\cup\mathcal{C}$. Make any identifications of the marked points implied by the edge identifications.

Lemma 3.1. $(X_t, \omega_t, \Sigma_t)$ converges to the surface $(X_0, \omega_0, \Sigma_0)$.

Proof. We provide a sketch only.

The edges of the boundary in the construction of $(X_0, \omega_0, \Sigma_0)$ give a finite union of parallel saddle connections on $(X_0, \omega_0, \Sigma_0)$. For each of these saddle connection, we can pick a long and thin embedded rectangle centered on this saddle connection.

The maps g_n can be defined to be an appropriate stretch in the orthogonal direction, defined on this rectangle minus the union of small neighborhoods about the zeros. See Figure 3.2. The maps g_n can be

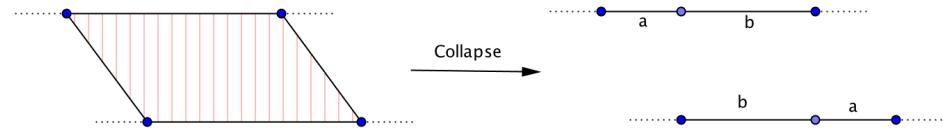


FIGURE 3.1. An example of a cylinder collapse. On the left, the two non-horizontal opposite edges are identified to give a horizontal cylinder. The surface is assumed to continue above and below, but only the cylinder is drawn. On the right, the cylinder is squashed to zero height, and becomes a pair of saddle connections.

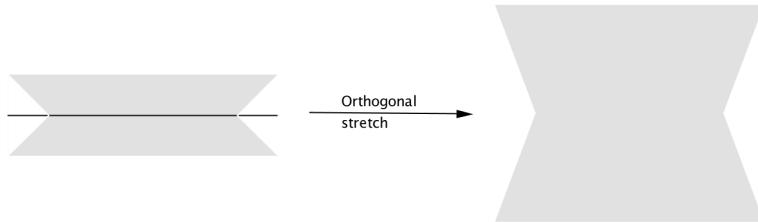


FIGURE 3.2. In this example, the horizontal line is a saddle connection in the image of \mathcal{C} under the collapse map. The map g_n takes a thin rectangle about this saddle connection and linearly stretches it in the orthogonal direction.

taken to be the identity outside the disjoint union of these rectangles; they are not defined on the neighborhoods about the zeros. \square

3.2. $\mathcal{H}(0,0)$ in the boundary of $\mathcal{H}(2)$. Figure 3.3 illustrates a sequence of surfaces in $\mathcal{H}(2)$ degenerating to $\mathcal{H}(0,0)$. It is clear that the degeneration cannot be achieved by surgeries supported near the zeros, since the ratio of the periods of α and β change as the surface degenerates, but α and β are supported away from the zeros.

3.3. Chaotic behavior in the vanishing part of the surface (McMullen). This example and the next give linear paths in period coordinates that converge in the sense above, but do not converge in Deligne-Mumford. These examples are helpful but not necessary to read the rest of the paper.

We are grateful to Curt McMullen for sharing with us this example, which shows how the chaos of the geodesic flow can be encoded in a convergent path which is even linear in period coordinates and consists entirely of unit area surfaces.

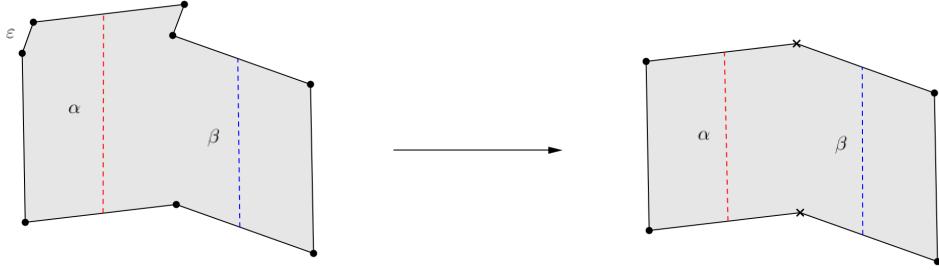


FIGURE 3.3. Opposite edges are identified. The edge labeled ϵ has length going to 0.

For $t > 0$ small and fixed y , let E_t be the torus $\mathbb{C}/\langle(0, 1), (1 - t, 0)\rangle$, and let F_t denote the torus $\mathbb{C}/\langle(0, 1), (t, y)\rangle$. Up to scaling, F_t is an orbit of geodesic flow, and it can be arranged¹ that $\{F_t : 0 < t < \epsilon\}$ is dense in \mathcal{M}_1 for each $\epsilon > 0$. Here \mathcal{M}_1 is the moduli space of complex tori, and we are forgetting the flat structure to get a point in \mathcal{M}_1 .

Let $(X_t, \omega_t) \in \mathcal{H}(1, 1)$ be F_t glued to E_t along a vertical slit of length t , as in Figure 3.4.

Observe that as $t \rightarrow 0$, the linear path (X_t, ω_t) accumulates in \mathcal{M}_2 on $E_0 \times \mathcal{M}_1$. Indeed, for any $F \in \mathcal{M}_1$, we may find t arbitrarily small so that F_t is arbitrarily close to F , and for these t we have that (X_t, ω_t) is close to F glued to E_0 at a point. (Note that, if we multiply F_t by $1/\sqrt{t}$ times the identify matrix, we obtain a unit area torus with a slit of length \sqrt{t} . Thus the slit is very small even compared to the decreasing “size” of F_t .)

3.4. Infinite death and rebirth of cylinders. The next example is not quite as chaotic as the previous one, but is more relevant in this paper. It shows the surprising phenomenon of “infinite death and rebirth” of cylinders in a *fixed homology class*, along a convergent linear path in period coordinates. This example could be easily modified to make the surfaces all of area exactly 1 (possibly changing the genus).

Consider the surfaces (X_t, ω_t) in $\mathcal{H}(2, 2)$ illustrated Figure 3.5. As $t \rightarrow 0$, these surfaces converge to two tori with one marked point each. When $t = 1/n$, there is a vertical cylinder, whose core curve is isotopic to $D_{\alpha_1}^n D_{\alpha_2}^{-n}(\gamma)$, where D_α denotes the Dehn twist about a curve α . By continuity, this cylinder also exists with t very close to $1/n$. But as t gets smaller or larger, it will disappear.

¹Indeed, $\mathbb{C}/\langle(0, 1), (t, y)\rangle$, viewed as a function of y for t fixed, is a horosphere for the geodesic flow on the bundle of Abelian differentials over \mathcal{M}_1 . Because geodesic flow is ergodic, almost any choice of y will work.

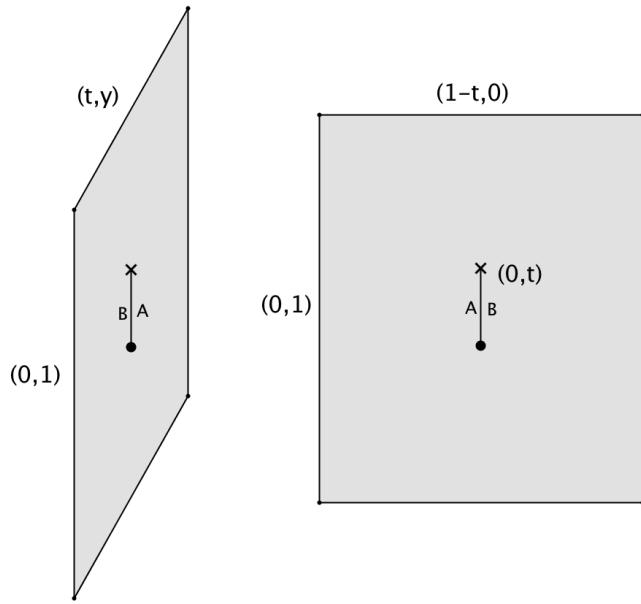


FIGURE 3.4. Opposite sides are identified, so each quadrilateral gives a torus. On the left is F_t and on the right is E_t . The edge identifications labeled A and B indicate that a slit in one torus is glued to a slit in the other.

Since α_1 and α_2 are equal in homology, we see that all the core curves of these cylinders are equal in homology.

In summary, infinitely many cylinders whose core curves are homologous to γ are “born” and “die” as $t \rightarrow 0$. Although the core curves of these cylinders are homologous, they are not isotopic.

4. CYLINDER TWISTS

In this section we recall some background, and finally prove a result (Theorem 4.7) that improves previous knowledge of cylinder deformations. The reader may find it helpful to become familiar with the Cylinder Deformation Theorem, for example reading at least the introduction of [Wri15a], before proceeding.

Note that when translation surfaces have marked points, if there is a zero inside of a cylinder, that cylinder is actually considered to be two adjacent cylinders. In other words, the interior of cylinders by definition may not contain marked points.

4.1. Cylinders and horocycle flow. Suppose (X, ω) is a translation surface, and C is a cylinder on (X, ω) with circumference c , height

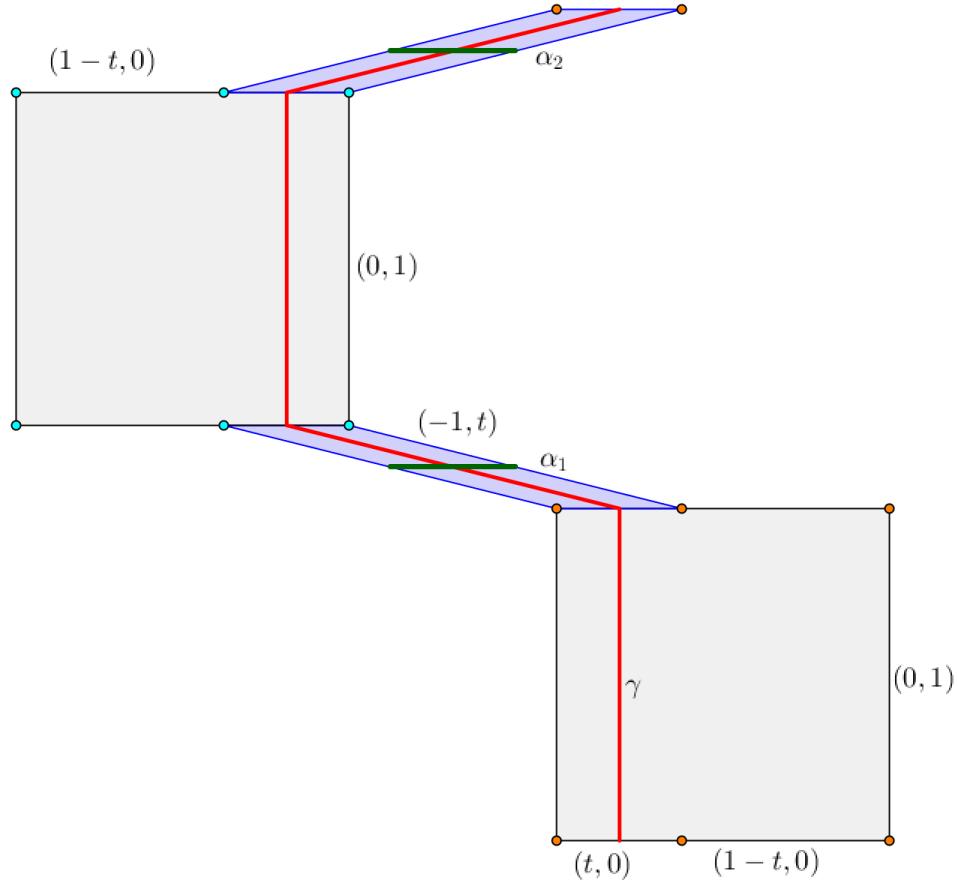


FIGURE 3.5. This figure shows two tori glued together with two cylinders. The tori are both 1 by 1, i.e., $\mathbb{C}/\mathbb{Z}[i]$. The cylinders have height t . The bottom cylinder has a saddle connection going from its bottom boundary to its top boundary with holonomy $(-1, t)$, and the top cylinder has a saddle connection going from its bottom boundary to its top boundary with holonomy $(1, t)$. The core curves of the bottom and top cylinders are α_1 and α_2 respectively.

h , modulus h/c , and core curve α . We assume an orientation of α is chosen. We will consider the cylinder twist $u_t^C(X, \omega)$, which is defined as follows. Rotate the surface by the appropriate angle $-\theta$ so that C becomes horizontal. Apply

$$u_t = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}$$

only to the cylinder C . This operation does not change the circumference of C or any of the saddle connections on the boundary of C . Then, rotate the surface by θ to return C to its original direction. The result is $u_t^C(X, \omega)$.

In other words, u_t^C shears the cylinder C , leaving the rest of the surface unchanged. The parameterization is such that it u_t^C has period c/h , i.e., when $t = c/h$ then $u_t^C(X, \omega) = (X, \omega)$.

There is a Poincare duality isomorphism

$$H_1(X - \Sigma, \mathbb{Z}) \simeq H_1(X, \Sigma, \mathbb{Z})^* = H^1(X, \Sigma, \mathbb{Z})$$

given by intersection number. Given a class $\alpha \in H_1(X - \Sigma, \mathbb{Z})$ we denote the dual class in $H^1(X, \Sigma, \mathbb{Z})$ by $I_\Sigma(\alpha)$. Note that there are natural maps

$$H_1(X - \Sigma, \mathbb{Z}) \rightarrow H_1(X, \mathbb{Z}) \quad \text{and} \quad p : H^1(X, \Sigma, \mathbb{Z}) \rightarrow H^1(X, \mathbb{Z}).$$

For a class $\alpha \in H_1(X, \mathbb{Z})$, we will denote its Poincare dual by $I(\alpha) \in H^1(X, \mathbb{Z})$.

We will consider the core curve α of C to be an element of both $H_1(X - \Sigma, \mathbb{Z})$ and $H_1(X, \mathbb{Z})$ without using different notation, and we have $p(I_\Sigma(\alpha)) = I(\alpha)$.

Define $u^C(X, \omega) = e^{i\theta} h I_\Sigma(\alpha)$, where θ is the argument of $\int_\alpha \omega$. We refer to θ as the direction of the cylinder. The derivative of the path $u_t^C(X, \omega)$ is equal to $u^C(X, \omega)$. Compare to Lemma 2.4 and Remark 2.5 in [Wri15a].

We will say that a collection of parallel saddle connections and cylinders is consistently oriented if each saddle connection, and each core curve of a cylinder, is given an orientation so that all their holonomies are positive multiples of each other. Whenever we speak of a collection of cylinders in this paper, we will assume it is consistently oriented.

Given a collection of cylinders $\mathcal{C} = \{C_1, \dots, C_k\}$, define $u_t^\mathcal{C}(X, \omega) = u_t^{C_1} \cdots u_t^{C_k}(X, \omega)$ to be the result of simultaneously shearing each C_i . Suppose C_i has height h_i , and define $u^\mathcal{C} = e^{i\theta} \sum h_i I_\Sigma(\alpha_i)$, where θ is the direction of the cylinders. The derivative of the path $u_t^\mathcal{C}(X, \omega)$ is $u^\mathcal{C}(X, \omega)$. Compare to Lemma 2.4 and Remark 2.5 in [Wri15a].

Let \mathcal{M} be an affine invariant submanifold. Recall that if $(X, \omega) \in \mathcal{M}$, then two cylinders on (X, ω) are said to be \mathcal{M} -parallel if they are parallel, and if they remain parallel on all sufficiently nearby surface in \mathcal{M} [Wri15a, Definition 4.6]. (More precisely: There is a small simply connected neighborhood of (X, ω) on which both cylinders persist, and whose intersection with \mathcal{M} is connected. We require the two cylinders to be parallel on all surfaces in this neighborhood intersect \mathcal{M} . See [LNW, Section 2.1] for the case when \mathcal{M} has self-crossings.)

If two cylinders are \mathcal{M} -parallel, then the ratio of their circumferences must remain constant on all nearby deformations [Wri15a, Lemma 4.7]. The following is [Wri15a, Theorem 5.1].

Theorem 4.1 (Cylinder Deformation Theorem). *Let \mathcal{M} be an affine invariant submanifold. If \mathcal{C} is an equivalence class of \mathcal{M} -parallel cylinders on $(X, \omega) \in \mathcal{M}$, then $u_t^{\mathcal{C}}(X, \omega)$ stays in \mathcal{M} for all $t \in \mathbb{R}$. Equivalently, $\sum h_i I_{\Sigma}(\alpha_i) \in T(\mathcal{M})$.*

The proof of Theorem 4.1 uses the following result.

Theorem 4.2 (Smillie-Weiss). *Every closed, horocycle flow invariant subset of a stratum contains a horizontally periodic surface.*

4.2. Cautionary examples. The possibility that small deformations can create new cylinders \mathcal{M} -parallel to the existing cylinders is a major complicating difficulty in this paper, so we encourage the reader to consider the following two examples in detail.

Example 4.3. When \mathcal{M} is a stratum, two cylinders are \mathcal{M} -parallel if and only if they are homologous. Take \mathcal{M} to be a stratum that contains a surface with a pair of homologous cylinders. Consider a surface in this stratum with a pair of homologous cylinders. Keeping the circumference constant, collapse one cylinder so that it has zero height, without degenerating the surface. Let the resulting surface be (X, ω) , and let \mathcal{C} be the one remaining cylinder from the homologous pair. Then, for many deformations, such as the reverse of the path we just described, a new cylinder will appear homologous to the one cylinder in \mathcal{C} .

Example 4.4. Take \mathcal{M} to be an eigenform locus in $\mathcal{H}(1, 1)$ [McM03], see also [Cal04]. Thus \mathcal{M} is rank 1, and so all parallel cylinders are automatically \mathcal{M} -parallel, see [Wri15b, Propositions 4.19, 4.20] or [Wri15a, Theorem 1.5]. Consider a surface $(X, \omega) \in \mathcal{M}$ with a two cylinder direction. Let \mathcal{C} be these two cylinders. Most small deformations of (X, ω) in \mathcal{M} result in a surface where there are three cylinders in the direction of the cylinders in \mathcal{C} .

Furthermore, McMullen has established the following result [McM15, McM14]. In each eigenform locus \mathcal{M} in genus 2, there are linear paths (X_t, ω_t) in period coordinates that converge as $t \rightarrow 0$ to a surface in the eigenform locus (not its boundary) with the following properties. For some $\varepsilon > 0$ and all $t \in (0, \varepsilon)$, the surfaces (X_t, ω_t) are all horizontally periodic. Furthermore, absolute periods are constant along the path (X_t, ω_t) , and as $t \rightarrow 0$, the circumference of the smallest horizontal cylinder goes to infinity.

Since absolute periods are constant along the path (X_t, ω_t) , the circumference of cylinders are locally constant. In order for the circumference of the smallest horizontal cylinder to go to infinity as $t \rightarrow 0$, the horizontal cylinders must disappear and be replaced with longer and longer horizontal cylinders. Infinitely many horizontal cylinders must appear and then die off. The “lifespan” of each horizontal cylinder goes to zero as $t \rightarrow 0$. Furthermore, since these \mathcal{M} are rank 1, all the core curves of the horizontal cylinders in this family are \mathcal{M} collinear.

This fact that \mathcal{M} -parallel cylinders in a given direction may appear and disappear infinitely many times along finite paths is a major source of complication in this paper.

See [HW] for related examples, with a very explicit analysis of how infinitely many cylinders appear and disappear along a small path (X_t, ω_t) , $0 < t < \varepsilon$. (In most of the examples in [HW], the cylinders are not \mathcal{M} -parallel.)

4.3. Non-standard twists. Given an equivalence class \mathcal{C} of \mathcal{M} -parallel cylinders with core curves α_i , define its *twist space* to be

$$(\text{span}_{\mathbb{C}} I_{\Sigma}(\alpha_i)) \cap T(\mathcal{M}).$$

This space describes all deformations that remain in \mathcal{M} that arise from deforming cylinders in \mathcal{C} .

We will call $u^{\mathcal{C}}$ the standard twist in \mathcal{C} . By the Cylinder Deformation Theorem, it is always contained in the twist space.

Warning 4.5. The definition of the standard twist depends on (X, ω) . For example, suppose $(X', \omega') \in \mathcal{M}$ is close enough to (X, ω) so that all cylinders in \mathcal{C} persist at (X', ω') . It may be that the heights of the cylinders in \mathcal{C} at (X, ω) and (X', ω') are not equal. Worse yet, it is possible that at (X', ω') there are cylinders \mathcal{M} -parallel to those in \mathcal{C} but not contained in \mathcal{C} . Either of these possibilities show that the standard twist in an equivalence class of \mathcal{M} -parallel cylinders is not locally constant, even when the cylinders persist.

Lemma 4.6. *Suppose that $(X, \omega) \in \mathcal{M}$. Let \mathcal{C} be the collection of all cylinders in (X, ω) in some direction, so in particular all cylinders in \mathcal{C} are parallel. Let $\mathcal{C} = \{C_1, \dots, C_k\}$ and suppose that C_i has circumference c_i , modulus m_i , height h_i , and core curve α_i . Then the subspace of*

$$\left\{ \sum t_i c_i I_{\Sigma}(\alpha_i) : (t_1, \dots, t_k) \in \mathbb{C}^k \right\}$$

that is contained in $T(\mathcal{M})$ is defined by linear equations on the variables t_i with coefficients in \mathbb{Q} that are satisfied when all $t_i = m_i$.

In this lemma, and for the remainder of the paper, we choose to parameterize cylinder deformations as linear combinations of $c_i I_\Sigma(\alpha_i)$ rather than $I_\Sigma(\alpha_i)$. Without this scaling of the coordinates, the rationality assertion (coefficients in \mathbb{Q}) of this lemma would be false.

Proof. This follows from a standard argument. Namely, the collection of all surfaces obtained by twisting the cylinders arbitrary amounts forms a torus of dimension k in the stratum. All twists that stay in \mathcal{M} give rise to a subtorus, which must be defined by rational equations in the coordinates. Since the $t_i = e^{i\theta} m_i$ is the standard twist and stays in \mathcal{M} by the Cylinder Deformation Theorem, all these rational equations must be satisfied at $t_i = m_i$. See [Wri15a, Corollary 3.4] for details. \square

Theorem 4.7. *The twist space of an equivalence class of \mathcal{M} -parallel cylinders is contained in the span of the standard twist and $\ker(p) \cap T(\mathcal{M})$.*

The proof will be by contradiction, using the following lemma.

Lemma 4.8. *Suppose that $(X, \omega) \in \mathcal{M}$ has an equivalence class of \mathcal{M} -parallel horizontal cylinders that supports two twists u, u' such that $p(u)$ and $p(u')$ are not collinear.*

Then there is an $(X', \omega') \in \mathcal{M}$ with the same property, such that (X', ω') is horizontally periodic and has a collection $\mathcal{C}_1, \dots, \mathcal{C}_r$ of equivalence classes of \mathcal{M} -parallel horizontal cylinders such that $p(u^{\mathcal{C}_i}), i = 1, \dots, r$ is a basis of a Lagrangian in $p(T(\mathcal{M}))$. We may also assume that u and u' are supported on \mathcal{C}_1 .

Here, by definition, r is the rank of \mathcal{M} .

Proof. This follows by a non-trivial argument that uses Theorem 4.2 repeatedly. The exact argument required appears in [Wri15a, Section 8]. \square

Proof of Proposition. Suppose in order to find a contradiction that (X, ω) has an equivalence class of \mathcal{M} -parallel horizontal cylinders that supports two twists u, u' such that $p(u)$ and $p(u')$ are not collinear. Let (X', ω') be as in the previous lemma.

Note that $p(u), p(u')$ together with $p(u^{\mathcal{C}_i}), i = 1, \dots, r$ span an isotropic subspace. Since $p(u^{\mathcal{C}_i}), i = 1, \dots, r$ are a basis for a Lagrangian, $p(u), p(u')$ can be extended to a basis for this Lagrangian by adding some of the $p(u^{\mathcal{C}_i})$. Thus, without loss of generality, we may assume that $p(u^{\mathcal{C}_i}), i = 1, \dots, r - 1$ together with u, u' span this Lagrangian.

This would imply that $p(u^{\mathcal{C}_r})$ can be written as a linear combination of $p(u), p(u')$ and $p(u^{\mathcal{C}_i}), i = 1, \dots, r - 1$ with *real* coefficients. (The

coefficients can be assumed to be real because all these vectors are contained in the real vector space $H^1(X, \Sigma, \mathbb{R})$.)

By definition, each $p(u^{\mathcal{C}_i})$ is dual to a real positive linear combination of core curves of cylinders in \mathcal{C}_i . Similarly, both u and u' are dual to a real linear combination of core curves of cylinders in \mathcal{C}_1 . Hence, we conclude that some positive linear combination of core curves of cylinders in \mathcal{C}_r is homologous to a real linear combination of core curves of the $\mathcal{C}_i, i = 1, \dots, r-1$.

However, the fact that the $p(u^{\mathcal{C}_i}), i = 1, \dots, r$ are linearly independent implies that the \mathcal{C}_r can be made not horizontal on a nearby surface $(X'', \omega'') \in \mathcal{M}$ while keeping the $\mathcal{C}_i, i = 1, \dots, r-1$ horizontal. This contradicts the homology relation in the previous paragraph. \square

5. CYLINDER FINITENESS

This section is the technical heart of this paper. Its purpose is to show the following.

Theorem 5.1 (Cylinder Finiteness Theorem). *Fix an affine invariant submanifold \mathcal{M} . There are two finite sets $S_1 \subset \mathbf{k}(\mathcal{M})$ and $S_2 \subset \mathbb{Q}$, such that if $(X, \omega) \in \mathcal{M}$, and \mathcal{C} is an equivalence class of \mathcal{M} -parallel cylinders on (X, ω) , then the ratio of circumferences of any two cylinders in \mathcal{C} is in S_1 .*

Furthermore, if C_i has circumference c_i and core curve α_i , then the twist space for \mathcal{C} at (X, ω) is a subspace of

$$\left\{ \sum t_i c_i I_\Sigma(\alpha_i) \right\}$$

defined by linear equations on the t_i with coefficients in S_2 that are satisfied when all $t_i = m_i$ are the moduli of the cylinders C_i .

Besides the statement of Theorem 5.1, no results from this section will be used in the rest of the paper. The strength of the statement is that S_1 and S_2 are finite. Indeed, it is easy to see that the ratio of two \mathcal{M} -parallel cylinders must be contained in $\mathbf{k}(\mathcal{M})$, see [Wri15a, Section 7], and we saw in Lemma 4.6 that twist spaces are always defined by linear equations on the t_i with rational coefficients.

Corollary 5.2. *There is a finite set $S'_2 \subset \mathbb{Q}$ depending only on S_2 and the maximum size of an equivalence class of \mathcal{M} -parallel cylinders such that the following is true.*

Suppose that \mathcal{C} is an equivalence class of \mathcal{M} -parallel cylinders on $(X, \omega) \in \mathcal{M}$, such that the ratio of moduli of cylinders in \mathcal{C} cannot be changed by deforming the surface in \mathcal{M} . Then the ratio of moduli of cylinders in \mathcal{C} is in S'_2 .

By Theorem 4.7, the assumption on \mathcal{C} is always satisfied if $\ker(p) \cap T(\mathcal{M}) = \{0\}$, and hence Corollary 1.6 is a special case of Corollary 5.2. A trivial bound for the maximum size of an equivalence class of \mathcal{M} -parallel cylinders is given by the maximum size of a set of parallel cylinders in the stratum.

Proof of Corollary. The twist space is defined by equations on the t_i with coefficients in S_2 , and the only solution is $t_i = m_i$. There are only finitely many such systems of equations, and we may take S'_2 to be the set of all ratios of t_i for the solutions of all these systems. \square

The strategy of the proof of Theorem 5.1 is to use the following corollary of the quantitative recurrence results of Minsky-Weiss [MW02]. For any affine invariant submanifold \mathcal{M} , let \mathcal{H}^1 denote the set of unit area surfaces in \mathcal{M} .

Theorem 5.3 (Minsky-Weiss). *Fix a stratum \mathcal{H} . For every $\varepsilon > 0$, there is a compact subset $K \subset \mathcal{H}^1$, such that if (X, ω) has no horizontal saddle connections of length less than ε , then the horocycle flow orbit of (X, ω) intersects K .*

We wish to use this result to show that any equivalence class of \mathcal{M} -parallel cylinders can be “moved” to a fixed compact set in a way that the smallest cylinder becomes circumference 1, and then use a compactness argument to obtain finiteness. Both parts of this strategy encounter major technical difficulties. These are handled in the next two subsections, after which Theorem 5.1 follows easily using Theorem 5.3.

We will make repeated use of the following matrices:

$$u_t = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad g_t = \begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}.$$

5.1. Saddle connections parallel to cylinders. Recall that a cylinder is said to be \mathcal{M} -parallel to a saddle connection if, locally on \mathcal{M} , the holonomy of the saddle connection is a fixed real multiple of the holonomy of the core curve of the cylinder.

Proposition 5.4. *Fix an affine invariant submanifold \mathcal{M} . Then there is a positive constant R_{sc} such that the follow holds. For every $(X, \omega) \in \mathcal{M}$, and for every saddle connection w on (X, ω) , if there is a cylinder \mathcal{M} -parallel to w , then there is cylinder \mathcal{M} -parallel to w whose circumference is at most R_{sc} times the length of w and whose area is at least $1/R_{sc}$ times the sum of the areas of the cylinders \mathcal{M} -parallel to w .*

In the next lemma, we will consider positive linear combinations of parallel consistently oriented saddle connections. Recall that consistently oriented means that the holonomies are positive multiples of each other. The key property that such linear combinations have in common with individual saddle connections is that if the holonomy does not change under a standard cylinder deformation, then the saddle connections are disjoint from the cylinders. We consider a saddle connection and a cylinder to be disjoint if the interior of a saddle connection is disjoint from the interior of the cylinder, so for example a saddle connection on the boundary of a cylinder is considered to be disjoint from the cylinder.

Recall that two relative homology classes γ and w are called \mathcal{M} -collinear if, for some $c \in \mathbb{R}$, the equation $\int_{\gamma} \omega = c \int_w \omega$ is one of the linear equations locally defining \mathcal{M} in period coordinates. If, as in the next lemma, γ is an absolute homology class but w is not, this means that w will behave like an absolute homology class for surfaces in \mathcal{M} .

Lemma 5.5. *Suppose that \mathcal{C} is an equivalence class of \mathcal{M} -parallel cylinders on $(X, \omega) \in \mathcal{M}$. Let w be a linear combination of saddle connections on (X, ω) that are disjoint from the cylinders in \mathcal{C} , and assume that w is \mathcal{M} -collinear to an absolute homology class γ .*

Then for $(X', \omega') \in \mathcal{M}$ sufficiently close to (X, ω) , if \mathcal{C}' is the equivalence class of cylinders \mathcal{M} -parallel to cylinders in \mathcal{C} , and w' is a positive linear combination of consistently oriented parallel saddle connections on (X', ω') equal to w in relative homology, then the cylinders in \mathcal{C}' are disjoint from w' on (X', ω') .

Let \mathcal{U} be a neighborhood of (X, ω) in \mathcal{M} , small enough so that all cylinders in \mathcal{C} , persist on all surfaces in \mathcal{U} . In the lemma, “close enough” means $(X', \omega') \in \mathcal{U}$. Note that \mathcal{C}' may be strictly larger than \mathcal{C} .

Proof. We assume γ is chosen so that $\int_{\gamma} \omega = \int_w \omega$.

Let $u = u^{\mathcal{C}}(X, \omega)$ be the standard cylinder twist in the cylinders of \mathcal{C} at (X, ω) , and let $u' = u^{\mathcal{C}'}(X', \omega')$ be the standard twist in the cylinders in \mathcal{C}' at (X', ω') .

By assumption, $u(w) = 0$. Theorem 4.7 gives that $p(u)$ and $p(u')$ are collinear. Hence $u'(w) = u'(\gamma)$ is proportional to $u(\gamma) = u(w) = 0$, so we get that $u'(w) = u'(w') = 0$. The result follows. \square

Lemma 5.6. *Fix an affine invariant submanifold \mathcal{M} , and suppose that $(X, \omega) \in \mathcal{M}^1$ and that w is a positive linear combination of consistently oriented horizontal saddle connections on (X, ω) . Suppose that w is \mathcal{M} -collinear to an absolute homology class.*

Then there is a constant $\delta > 0$ depending only on the stratum such that there is an equivalence class \mathcal{C} of \mathcal{M} -parallel cylinders on (X, ω) that is disjoint from w and that has area at least δ .

Proof. Set δ to be 1 divided by twice the maximum number of parallel cylinders on a surface in the given stratum. Let $(X', \omega') \in \mathcal{M}$ be a horizontally periodic surface in the horocycle flow orbit closure of (X, ω) .

Let \mathcal{C} be an equivalence class of \mathcal{M} -parallel horizontal cylinders on (X', ω') that occupy at least 2δ of the area.

Pick t so that $u_t(X, \omega)$ is close enough to (X', ω') so all cylinders in \mathcal{C} persist at $u_t(X, \omega)$, and so that they occupy at least δ of the area.

Let \mathcal{C}' be the equivalence class of cylinders on $u_t(X, \omega)$ that are \mathcal{M} -parallel to cylinders in \mathcal{C} . By Lemma 5.5, the cylinders in \mathcal{C}' are disjoint from w .

Pulling the cylinders of \mathcal{C}' back to (X, ω) using u_{-t} gives the result. \square

Proof of Proposition 5.4. Assume to the contrary that there is a sequence $(X_n, \omega_n) \in \mathcal{M}^1$ of unit area surfaces where there is a saddle connection w_n that is \mathcal{M} -parallel to a cylinder, but that the set \mathcal{C}_n of such cylinders can be broken up into two disjoint subsets $\mathcal{C}_n = \mathcal{C}'_n \cup \mathcal{C}''_n$ such that all cylinders in \mathcal{C}'_n have circumference, divided by the length of w_n , going to infinity uniformly in n , and the cylinders in \mathcal{C}''_n have area, divided by the total area of \mathcal{C}_n , going to zero uniformly in n .

Slightly deforming the surfaces, we may assume that w_n is not parallel to any saddle connection that it is not \mathcal{M} -parallel to. (Any cylinders created in this process that are \mathcal{M} -parallel to those in \mathcal{C}_n can be assumed to have arbitrarily small area.) Without loss of generality, we will assume that any saddle connection parallel to w_n is at least as long as w_n .

Rotating the surfaces, we may assume that all w_n are horizontal. The reason we choose to have w_n horizontal is so that w_n and \mathcal{C}_n will be preserved by the horocycle flow. Applying g_t , we may assume that w_n has length 1. By applying the cylinder stretch in \mathcal{C}_n (the deformation with derivative $iu^{\mathcal{C}_n}$, which increases heights of cylinders) we can increase the ratio of the area of the cylinders in \mathcal{C}_n to the total area. This changes the area of the surface, but we can then apply matrices of the form

$$\begin{pmatrix} 1 & 0 \\ 0 & a \end{pmatrix}$$

to renormalize the area. The result is that we may assume that the area of \mathcal{C}_n goes to 1 as n goes to infinity. Hence the area of \mathcal{C}'_n goes to 1 as n goes to infinity.

By applying horocycle flow and using Theorem 5.3, we may assume that all the (X_n, ω_n) lie in some fixed compact set. Hence, by passing to a subsequence, we may assume that they converge to a limit in the same stratum. Say the limit is (X, ω) .

This limit has a finite union of horizontal saddle connections of length 1, which we will call w . (Multiplicity is allowed.) This is the limit of w_n on (X_n, ω_n) , and has the same relative homology class as the w_n when n is large enough. The relative homology class of w is, by assumption, \mathcal{M} -collinear to any γ which is the core curve of a cylinder in \mathcal{C}_n on (X_n, ω_n) .

Hence by Lemma 5.6, we get the existence of an equivalence class of \mathcal{M} -parallel cylinders \mathcal{D} on (X, ω) that is disjoint from w , and occupies area at least δ . Pick n large enough so that \mathcal{C}'_n has area greater than $1 - \delta/2$, and such that the cylinders of \mathcal{D} persist on (X_n, ω_n) and occupy area at least $\delta/2$.

Let \mathcal{D}_n be the set of cylinders on (X_n, ω_n) that are \mathcal{M} -parallel to the cylinders in \mathcal{D} . Lemma 5.5 shows that \mathcal{D}_n is disjoint from w_n . Furthermore, \mathcal{D}_n contains cylinders whose circumference and area are uniformly bounded above and below, namely those coming from \mathcal{D} . Hence \mathcal{D}_n is not contained in \mathcal{C}_n , and hence no cylinder in \mathcal{D}_n is contained in \mathcal{C}_n and vice versa. (Since they are equivalence classes, if \mathcal{C}_n and \mathcal{D}_n had a single cylinder in common, they would be equal.)

However, by area considerations, some cylinder in \mathcal{D}_n must intersect some cylinder in \mathcal{C}_n . This is a contradiction, because the cylinder deformation in the \mathcal{D}_n changes the circumferences of at least one cylinder in \mathcal{C}_n , without changing the holonomy of w_n , contradicting the fact that w_n and the cylinders of \mathcal{C}_n are \mathcal{M} -parallel. \square

5.2. Bounding cylinder ratios.

Proposition 5.7. *Fix an affine invariant submanifold \mathcal{M} . Then there is a constant R_{cyl} such that any ratio of circumferences of any pair of \mathcal{M} -parallel cylinders on any surface in \mathcal{M} is at most R_{cyl} .*

The following basic lemma will be helpful in combination with Theorem 4.7 to understand how new cylinders can “appear” in equivalence classes as the surface is deformed.

Lemma 5.8. *Suppose that $\alpha_i, i = 1, \dots, n$ and $\beta_j, j = 1, \dots, m$ are all core curves of a set of consistently oriented distinct parallel cylinders*

on a translation surface (X, ω) . Suppose furthermore that in homology there is some linear relation

$$\sum a_i \alpha_i = \sum a'_i \alpha_i + \sum b_j \beta_j$$

with all a_i and a'_i and b_j positive real numbers.

Then for each j_0 there is a subset $J \subset \{1, \dots, m\}$ containing j_0 and two subsets $I, I' \subset \{1, \dots, n\}$ such that in homology,

$$\sum_{i \in I} \alpha_i = \sum_{i \in I'} \alpha_i + \sum_{j \in J} \beta_j.$$

In particular, for each j we have

$$\left| \int_{\beta_j} \omega \right| \leq \sum_{i=1}^n \left| \int_{\alpha_i} \omega \right|.$$

Recall that consistently oriented means that the $\int_{\alpha_i} \omega, \int_{\beta_j} \omega$ are all positive multiples of each other.

Proof. Without loss of generality, assume the cylinders are horizontal. Cut out all the α_i and β_j from X , to obtain a finite union X_1, \dots, X_k of connected components, each of which carries a flat structure and has horizontal boundary.

Suppose that β_{j_0} is part of the top boundary of X_1 . Iteratively perform the following operation, illustrated in Figure 5.1. Start with X_1 . If there is a bottom boundary component that is a β_j , glue on the piece X_p that contains this β_j in the top. In the object thus obtained, if there is a bottom boundary component that is a different β_j , glue on the piece $X_{p'}$ that contains this β_j in the top.

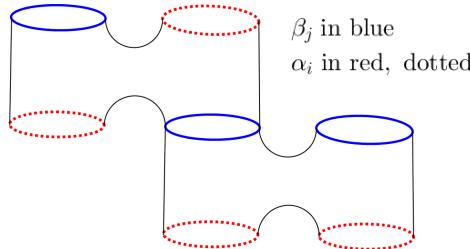


FIGURE 5.1. Proof of Lemma 5.8.

We claim that when this process is completed, β_{j_0} is contained in the top boundary of the resulting object (and not in the bottom—there are no β_j in the bottom boundary). Indeed, otherwise we could find a curve γ going through the pieces and crossing only the β_j curves and

not crossing any α_i curves. This curve would contradict the assumed linear relation.

The lemma is now proved, with J the set of j such that β_j occurs in the top boundary, I the set of i such that α_i occurs in the bottom boundary, and I' the set of i such that α_i occurs in the top boundary. \square

Lemma 5.9. *Let \mathcal{M} be an affine invariant submanifold, and $(X, \omega) \in \mathcal{M}$. Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be an equivalence class of \mathcal{M} -parallel cylinders on (X, ω) . Suppose that $(X', \omega') \in \mathcal{M}$ is a small deformation of (X, ω) , so all the C_i persist on (X', ω') . Suppose that C_0 is a cylinder on (X', ω') that is \mathcal{M} -parallel to the $C_i, i = 1, \dots, n$. Then the circumference of C_0 is at most the sum of the circumferences of the core curves of the cylinders in \mathcal{C} .*

Example 5.10. It is however possible that the core curve of C_0 is not in the span of the core curves of the $C_i, i = 1, \dots, n$. In Figure 5.2 we draw an example of a locus \mathcal{M} of degree 2 torus covers in $\mathcal{H}(1, 1)$. The surface on the right is the (X, ω) . Many deformations, such as those on the left, have two “new” cylinders that are \mathcal{M} -parallel to the cylinder that persists from (X, ω) . However, these cylinders are not homologous to the cylinder that persists from (X, ω) .

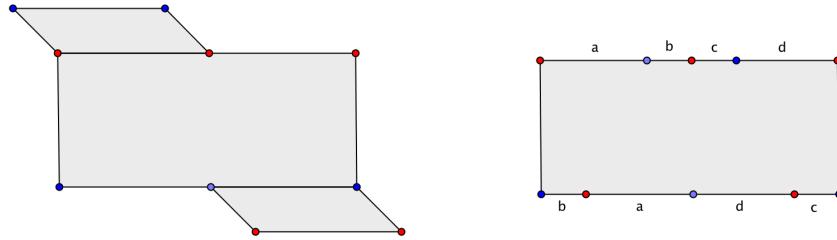


FIGURE 5.2

Proof. Let u be the standard twist in \mathcal{C} at (X, ω) . Let \mathcal{C}' be the equivalence class of C_1 at (X', ω') , so $\mathcal{C} \subset \mathcal{C}'$. At (X', ω') , the relative cohomology class u still describes a twist in the cylinders of \mathcal{C}' . This twist involves only the cylinders in \mathcal{C} , and not those in $\mathcal{C}' \setminus \mathcal{C}$.

Let u' be the standard twist at (X', ω') in the cylinders of \mathcal{C}' . By Theorem 4.7, by rescaling u' if necessary, we may assume $p(u) = p(u')$.

Suppose the core curves of the cylinders in \mathcal{C} are α_i , and the core curves of the cylinders of $\mathcal{C}' \setminus \mathcal{C}$ are β_i . Assume for convenience that the cylinders under consideration are all horizontal. We have

$$u = \sum h_i I_\Sigma(\alpha_i) \quad \text{and} \quad u' = \sum h'_i I_\Sigma(\alpha_i) + \sum \delta_i I_\Sigma(\beta_i),$$

where the δ_i are the heights of the new cylinders, and the h_i and h'_i are the heights of the cylinders in \mathcal{C} at (X, ω) and (X', ω') respectively. Since these are equal in absolute homology, we get

$$\sum h_i \alpha_i = \sum h'_i \alpha_i + \sum \delta_j \beta_j$$

in absolute homology. Lemma 5.8 now gives the result. \square

Proof of Proposition 5.7. Suppose to the contrary that there is a sequence of surfaces $(X_n, \omega_n) \in \mathcal{M}^1$ that has an equivalence class \mathcal{C}_n of \mathcal{M} -parallel cylinders where the ratio of the largest to the smallest circumference goes to infinity. Without loss of generality, assume that all these \mathcal{C}_n are horizontal, and that there are no horizontal saddle connections on (X_n, ω_n) that are not \mathcal{M} -parallel to the cylinders in \mathcal{C}_n . (If there are horizontal saddle connections on (X_n, ω_n) that are not \mathcal{M} -parallel to the cylinders in \mathcal{C}_n , the surface can be deformed slightly within \mathcal{M} to remove this unnecessary coincidence.)

Let w_n be the shortest horizontal saddle connection on (X_n, ω_n) . By Proposition 5.4, there is a constant R_{sc} such that there is a cylinder in \mathcal{C}_n that has circumference at most R_{sc} times the length of w_n , and that has area at least $1/R_{sc}$ times the total area of \mathcal{C}_n . As in the proof of Proposition 5.4, we can assume that this cylinder has circumference 1 and area at least $1/(2R_{sc})$. Since the smallest horizontal saddle connection on (X_n, ω_n) has length at least $1/R_{sc}$, using Theorem 5.3 we may assume that (X_n, ω_n) converges to some surface (X, ω) in the same stratum.

There is at least one cylinder C on (X, ω) that is a limit of horizontal cylinders on (X_n, ω_n) of circumference 1 and area at least $1/(2R_{sc})$. Let \mathcal{C} be the equivalence class of C .

Let Q be the sum of the circumferences of the cylinders in \mathcal{C} . We now claim that for large n , all cylinders in \mathcal{C}_n are size at most $Q + 1$, which contradicts our original supposition that the ratio of the largest and smallest circumferences of cylinders in \mathcal{C}_n goes to infinity.

The claim follows directly from Lemma 5.9. \square

5.3. Proof the Cylinder Finiteness Theorem.

Lemma 5.11. *For any simply connected compact set K of a stratum, and any constant $L > 0$, the set of relative homology classes represented by a saddle connection of length at most L on some $(X, \omega) \in K$ is finite.*

The fact that K is simply connected means that the relative homology of different surfaces in X can be canonically identified. The proof is standard, but we give a sketch.

Proof. There is some $c > 1$ depending on K so that if a relative cohomology class is represented by a saddle connection of length L somewhere in K , then it is represented by a union of saddle connections of total length at most cL everywhere on K .

If there were infinitely many relative cohomology classes represented somewhere on K by saddle connections of length at most L , then on every $(X, \omega) \in K$ all of these relative cohomology classes would be represented by unions of saddle connections of total length at most cL . This contradicts the fact that there are at most finitely many saddle connections on (X, ω) of length at most cL . \square

Lemma 5.12. *Let \mathcal{U} be a set of translation surfaces with compact closure in the stratum. Let $L > 0$ be a large constant. Then \mathcal{U} is contained in a finite union of connected subsets \mathcal{V} , such that the following properties hold.*

- Suppose $(X, \omega) \in \mathcal{V}$ has a saddle connection of length at most L . Then this saddle connection exists everywhere in \mathcal{V} .
- Suppose $(X, \omega) \in \mathcal{V}$ has a pair saddle connection of length at most L that are parallel. Then these saddle connections remain parallel on all of \mathcal{V} .

Proof. It suffices to show this for any small set \mathcal{U} that is a round ball in some choice of local period coordinates.

Let $P \subset H_1(X, \Sigma, \mathbb{Z})$ be the set of relative homology classes that are represented by saddle connections of length at most cL at some surface in \mathcal{U} , where c is the constant from the proof of Lemma 5.11. This set is finite by Lemma 5.11. For each $\alpha, \beta \in P$, consider the set $H_{\alpha, \beta} \subset \mathcal{U}$ where the holonomy of α and β are parallel. In local coordinates, this set is a hypersurface defined by a single quadratic equation.

For each subset of pairs $J \subset P \times P$, consider the subset

$$\mathcal{V}_J = \bigcap_{(\alpha, \beta) \in J} H_{\alpha, \beta} \setminus \bigcup_{(\alpha, \beta) \notin J} H_{\alpha, \beta}$$

of surface in \mathcal{U} that are in $H_{\alpha, \beta}$ if and only if $(\alpha, \beta) \in J$. Each \mathcal{V}_J intersects \mathcal{U} in at most finitely many connected components. (Indeed, these intersections are semi-algebraic sets, which always have at most finitely many connected components [BCR98, Theorem 2.4.5].)

It is clear that the connected components of \mathcal{V}_J satisfy the two desired properties, since when a saddle connection is created or destroyed a pair of previously parallel saddle connections become not parallel or vice versa. \square

Corollary 5.13. *Let \mathcal{M} be an affine invariant submanifold and let $L' > 0$. Then each compact subset $K \subset \mathcal{M}$ can be covered by finitely*

many connected sets \mathcal{V} , such that the following property holds. If $(X, \omega) \in \mathcal{V}$ and C is a cylinder of length at most L' on (X, ω) , then the set of all cylinders \mathcal{M} -parallel to C is constant on \mathcal{V} , and the set of boundary saddle connections of these cylinders is also constant on \mathcal{V} .

Proof. This follows from Lemma 5.12 with constant L equal to L' times any upper bound R_{cyl} for the ratio of two \mathcal{M} -parallel cylinders in \mathcal{M} , which exists by Proposition 5.7. \square

Finally, we can conclude the proof of Theorem 5.1.

Proof of Theorem 5.1. Pick $\varepsilon = 1/R_{sc}$, where R_{sc} is the constant from Proposition 5.4. Let $K \subset \mathcal{H}^1$ be the compact set given by Theorem 5.3 such that every horocycle orbit of a unit area surface without saddle connections of length less than ε intersects K .

Cover K by finitely many connected sets \mathcal{V} as in the previous corollary, using $L' = 1$. For each \mathcal{V} and each cylinder C of circumference at most 1 on a surface in \mathcal{V} , this cylinder and the equivalence class \mathcal{C} of cylinders \mathcal{M} -parallel to C persist everywhere on \mathcal{V} , and the ratio of circumferences of cylinders in \mathcal{C} is constant on \mathcal{V} . There are only finitely many such C and \mathcal{C} for each \mathcal{V} .

Suppose that the cylinders in \mathcal{C} have core curves α_i and circumferences c_i . Each twist in this equivalence class can be written as

$$\sum t_i c_i I_\Sigma(\alpha_i)$$

for some t_i in \mathbb{C} . The t_i are coordinates for the set of all twists supported on \mathcal{C} . Note the circumferences c_i are not constant on \mathcal{V} , but their ratios are. The twist space of \mathcal{C} is also constant on \mathcal{V} .

Define the sets S_1 and S_2 as follows. For every \mathcal{V} , and every equivalence class \mathcal{C} of \mathcal{M} -parallel cylinders that contains a cylinder of circumference at most 1 somewhere on \mathcal{V} , add the ratios of the circumferences to S_1 . By Lemma 4.6, the twist space of \mathcal{C} can be defined by linear equations on the t_i with coefficients in \mathbb{Q} . Pick a finite set of such equations defining the twists space of \mathcal{C} , and add all the coefficients to S_2 . (The definition of S_2 is not canonical.)

We will prove the theorem with this choice of S_1 and S_2 . Take any surface $(X, \omega) \in \mathcal{M}^1$, and any equivalence class \mathcal{C} of \mathcal{M} -parallel cylinders on (X, ω) . Without loss of generality, assume these cylinders are horizontal.

If there are horizontal saddle connections not \mathcal{M} -parallel to the cylinders in \mathcal{C} , nudge (X, ω) to get (X', ω') where all cylinders in \mathcal{C} persist and stay horizontal, but where there are no such horizontal saddle connections accidentally parallel to the cylinders in \mathcal{C} . Let \mathcal{C}' be the set of

cylinders at (X', ω') that are \mathcal{M} -parallel to cylinders in \mathcal{C} , so $\mathcal{C} \subset \mathcal{C}'$. Note that the twist space for \mathcal{C} at (X, ω) embeds naturally in the twist space for \mathcal{C}' at (X', ω') via parallel transport. Furthermore, in the t_i coordinates, the twist space of \mathcal{C} is a coordinate subspace of the twist space for \mathcal{C}' , because twists in \mathcal{C}' are parallel transports of twists in \mathcal{C} if and only if they do not involve the cylinders of $\mathcal{C}' \setminus \mathcal{C}$, i.e. if the t_i corresponding to cylinders in $\mathcal{C}' \setminus \mathcal{C}$ are 0.

By Proposition 5.4, if we pick t such that the smallest circumference of a cylinder in \mathcal{C}' on $g_t(X', \omega')$ is exactly 1, then the shortest horizontal saddle connection is at most $1/R_{sc}$. By Theorem 5.3, we can find s so that $h_s g_t(X', \omega') \in K$. By the definition of S_1 and S_2 , this gives that the circumferences and twist space of \mathcal{C}' and hence also \mathcal{C} are described by the finite sets S_1 and S_2 as desired. \square

6. RECOGNIZABLE TWISTS

The goal of this section is to prove that certain *recognizable* twists span $T(\mathcal{M})$, and that a recognizable twist at a surface is still recognizable at nearby surfaces. This ability to robustly recognize the tangent space will allow us to compare the tangent space at a surface in \mathcal{M} to the tangent space at a surface in $\partial\mathcal{M}$ in the next section.

Remark 6.1. Before discussing what does work, it is helpful to mention what doesn't work. We could simply consider, for each surface (X, ω) , the span $E(X, \omega)$ of the standard twists in the full set of all cylinders in each direction. (Or, what is a bit better, we could also include twists that are in the closure of the one parameter families of deformations given by the standard twists. So, for example if a direction on (X, ω) has exactly two parallel cylinders, but their moduli aren't rational multiples of each other, this would include the twist in each of these cylinders separately.)

This $E(X, \omega)$ is defined without reference to any particular \mathcal{M} and is very poorly behaved. For example, square-tiled surfaces are dense in every stratum, and for each square-tiled surface $E(X, \omega)$ is just the tautological plane $\text{span}(\text{Re}(\omega), \text{Im}(\omega))$. In particular, $E(X, \omega)$ is nowhere continuous. Even worse, as far as we know it is possible for $E(X, \omega)$ to be very small even if the orbit closure of (X, ω) is large. For example, it is an open question whether $E(X, \omega) = \text{span}(\text{Re}(\omega), \text{Im}(\omega))$ implies that (X, ω) lies on a closed orbit [SW07, Question 5]. (For some progress on this question, see [LNW].)

6.1. Definitions and examples. We will begin in a fairly abstract setting, without reference to \mathcal{M} . The reader may consult Lemma 6.5 as they are reading the definitions for some motivation.

Let S_1, S_2 be two finite sets of real numbers. Suppose that (X, ω) is a translation surface, and let $\{C_i\}$ be the set of all cylinders in some direction on (X, ω) . Suppose C_i has height h_i , circumference c_i , modulus $m_i = h_i/c_i$, and core curve α_i .

We will say that the twist

$$u = \sum t_i c_i I_\Sigma(\alpha_i)$$

is (S_1, S_2) -recognizable if, for all subsets of the cylinders C_i with all ratios of circumferences in S_1 , the corresponding t_i satisfy all linear equations with coefficients in S_2 that the m_i do.

More precisely, we require that for any subset $\{C_i : i \in \mathcal{I}\}$ of the cylinders with $c_i/c_j \in S_1$ for all $i, j \in \mathcal{I}$, and any choice of $f_i \in S_2$ for each $i \in \mathcal{I}$,

$$\sum_{\mathcal{I}} f_i m_i = 0 \quad \text{implies} \quad \sum_{\mathcal{I}} f_i t_i = 0.$$

Note that the $c_i = c_i(X, \omega)$ and $m_i = m_i(X, \omega)$ vary with (X, ω) , and so a twist can be (S_1, S_2) -recognizable at a surface but not at some nearby surface.

Example 6.2. Suppose that in some direction (X, ω) has three cylinders, with circumferences $c_1 = 1, c_2 = 1, c_3 = 40$ and moduli $m_1 = 1, m_2 = 2, m_3 = 1$. Suppose $S_1 = S_2 = \{1\}$. Then the twist described by $t_1 = 1, t_2 = -3, t_3 = 0$ is (S_1, S_2) -recognizable. To check this, note that the only set of cylinders with all ratios of circumferences in S_1 is $\{C_1, C_2\}$. So we need to check that t_1, t_2 satisfy all linear equations with coefficients in S_2 that m_1, m_2 satisfy. In this case m_1, m_2 don't satisfy any linear equations with coefficients S_2 .

Next consider the same example but with $S_1 = \{1, 40, 1/40\}$ and $S_2 = \{1\}$. The full set of all three cylinders has all ratios of circumferences in S_1 . So we need to check that t_1, t_2, t_3 satisfy all linear equations with coefficients in S_2 that the m_1, m_2, m_3 satisfy. In this case the equation $m_1 = m_3$ holds, but $t_1 \neq t_3$, so the twist is not recognizable.

Next consider the same example but with $S_1 = \{1\}, S_2 = \{1, 2\}$. The only set of cylinders with all ratios of circumferences in S_1 is $\{C_1, C_2\}$, but now the equation $m_2 = 2m_1$ shows the twist is not (S_1, S_2) -recognizable.

Example 6.3. If \mathcal{C} is the collection of all cylinders in some direction on (X, ω) , then the standard twist in \mathcal{C} is (S_1, S_2) -recognizable for any S_1, S_2 . This is because this twist corresponds to $t_i = m_i$.

Remark 6.4. For many surfaces the standard twists will be the only (S_1, S_2) -recognizable twists. However at special surfaces, for example when there are extra parallel cylinders in some direction, there may be (S_1, S_2) -recognizable twists that are not standard.

6.2. Key results.

Lemma 6.5. *For any affine invariant submanifold \mathcal{M} , there are finite sets S_1, S_2 so that any (S_1, S_2) -recognizable twist at a surface in \mathcal{M} remains in \mathcal{M} .*

Pairs of sets (S_1, S_2) satisfying the conclusion of the lemma will be called \mathcal{M} -adapted. Note that enlarging an adapted set always gives another adapted set.

Proof. Take the S_1, S_2 given by the Cylinder Finiteness Theorem.

Suppose that u is a (S_1, S_2) -recognizable twist. It suffices to assume it is supported on a equivalence class of \mathcal{M} -parallel cylinders, since u is the sum of its restrictions to such equivalence classes.

Write $u = \sum t_i c_i I_\Sigma(\alpha_i)$. By assumption, the twist space of the equivalence class of \mathcal{M} -parallel cylinders can be defined by linear equations in the t_i with coefficients in S_2 . The twist space contains the standard twist given by $t_i = m_i$, so hence all of the equations defining the twist space must have $t_i = m_i$ as a solution.

By the definition of (S_1, S_2) -recognizable, the t_i defining u satisfy all such linear equations. \square

The key technical property of recognizable twists is given by the next lemma.

Lemma 6.6. *Suppose that u is a (S_1, S_2) -recognizable twist at (X, ω) . Then there is a neighborhood \mathcal{V} of (X, ω) such that for all $(X', \omega') \in \mathcal{V}$, the parallel translate of u to (X', ω') is in the span of the (S_1, S_2) -recognizable twists at (X', ω') .*

By span we always mean span over the complex numbers.

Proof. Let $C_i, i = 1, \dots, k$ be the collection of all cylinders in the direction of the cylinder twist u . Say

$$u = \sum_{i=1}^k t_i c_i(X, \omega) I_\Sigma(\alpha_i).$$

Let n be the maximum number of parallel cylinders on a surface in the stratum of (X, ω) . Consider

$$(m_1(X, \omega), \dots, m_k(X, \omega), 0, \dots, 0) \in \mathbb{R}^{k+n}.$$

Find a neighborhood $\mathcal{U} \subset \mathbb{R}^{k+n}$ of this point such that for all $(m'_1, \dots, m'_{k+n}) \in \mathcal{U}$, the m'_i do not satisfy any linear equations with coefficients in S_2 except those of the form $E_1 + E_2 = 0$, where E_1 is a linear equation with coefficients in S_2 satisfied by $(m_1(X, \omega), \dots, m_k(X, \omega))$, and E_2 is any linear equation in the remaining n variables with coefficients in S_2 .

We may find a neighborhood $\mathcal{U}_0 \times B_\varepsilon \subset \mathcal{U}$, where \mathcal{U}_0 is a neighborhood of $(m_1(X, \omega), \dots, m_k(X, \omega))$ in \mathbb{R}^k , and B_ε is the ball of radius ε in the remaining coordinates \mathbb{R}^n .

Define \mathcal{V} to be the set of (X', ω') sufficiently close to (X, ω) so that the following three conditions hold.

- (1) All the C_i persist at (X', ω') and the moduli $m_i(X', \omega')$ of the C_i at (X', ω') have $(m_1(X', \omega'), \dots, m_k(X', \omega')) \in \mathcal{U}_0$.
- (2) If C is a cylinder parallel to C_i at (X', ω') , and C is not one of the C_j , and there is some j so that the ratio of the circumference of C and C_j is in S_1 , then the modulus of C is less than ϵ .
- (3) If C_i and C_j have ratio of circumferences in S_1 at (X', ω') , then they have the same ratio of circumferences at (X, ω) .

Consider the equivalence relation on the C_i , generated by $C_i \sim C_j$ if C_i and C_j remain parallel on (X', ω') and the ratio of their circumferences at (X', ω') is in S_1 . Let $C_i, i \in \mathcal{I}$ be an equivalence class for this equivalence relation.

Claim: For each such equivalence class \mathcal{I} , the twist

$$u_{\mathcal{I}} = \sum_{i \in \mathcal{I}} t_i c_i(X, \omega) I_{\Sigma}(\alpha_i)$$

is a recognizable twist at (X', ω') .

Conclusion of the proof of lemma given the claim: u is the sum of the twists $u_{\mathcal{I}}$ over the different equivalence classes \mathcal{I} . Since each $u_{\mathcal{I}}$ is (S_1, S_2) -recognizable at (X', ω') , we see that u is in the span of the recognizable twists at (X', ω') .

Proof of claim: Because the ratios of circumferences of the $C_i, i \in \mathcal{I}$ have not changed in moving from (X, ω) to (X', ω') , the cohomology class $u_{\mathcal{I}}$ is proportional to

$$\sum_{i \in \mathcal{I}} t_i c_i(X', \omega') I_{\Sigma}(\alpha_i).$$

Suppose, in order to find a contradiction, that this twist is not recognizable at (X', ω') . By definition, there must be

- a collection of cylinders on (X', ω') parallel to the $C_i, i \in \mathcal{I}$ for which all ratios of circumferences are in S_1 , and

- a linear equation E with coefficients in S_2 that is satisfied by the moduli of these cylinders but not the t_i .

The t_i value for any cylinder other than the $C_i, i \in \mathcal{I}$ is 0, because only the cylinders $C_i, i \in \mathcal{I}$ appear in the definition of $u_{\mathcal{I}}$.

Let $D_j, j \in \mathcal{J}$ be the set of cylinders on (X', ω') that are parallel to some $C_i, i \in \mathcal{I}$ and have ratio of circumference with this C_i in S_1 , but which are not equal to any of the C_i (the full set of parallel cylinders considered on (X, ω)). By part (2) of the definition of \mathcal{V} , all D_j have modulus less than ε . By part (3), the set of $m_i(X', \omega'), i \in \mathcal{I}$ and $m_j(X', \omega'), j \in \mathcal{J}$ do not satisfy any linear equations with coefficients in S_2 , except those of the form $E_1 + E_2$, where E_1 is an equation on the $t_i, i \in \mathcal{I}$ that holds for the $m_i(X, \omega)$, and E_2 is any equation with coefficients in S_2 on the new t -variables corresponding to the $D_j, j \in \mathcal{J}$. So we may assume $E = E_1 + E_2$ has this form.

The $t_i, i \in \mathcal{I}$ satisfy E_1 by construction, and $u_{\mathcal{I}}$ has 0 value for the new t -variables corresponding to the $D_j, j \in \mathcal{J}$. Hence E is satisfied by the t_i , and we get a contradiction. We conclude that $u_{\mathcal{I}}$ is in fact recognizable at (X', ω') . \square

Lemma 6.7. *Let (S_1, S_2) be arbitrary. There is an open, $GL(2, \mathbb{R})$ invariant set in \mathcal{M} on which the span of the (S_1, S_2) -recognizable twists is locally constant.*

Proof. Consider the open, $GL(2, \mathbb{R})$ invariant set where the span of the (S_1, S_2) -recognizable twists has maximal dimension. On this locus, the span is locally constant by Lemma 6.6. \square

Proposition 6.8. *Suppose (S_1, S_2) is \mathcal{M} -adapted. There is a finite union of proper affine invariant submanifolds of \mathcal{M} such that if $(X, \omega) \in \mathcal{M}$ is not contained in this finite union then the (S_1, S_2) -recognizable twists at (X, ω) generate the tangent space to \mathcal{M} at (X, ω) .*

Proof. Consider the open $GL(2, \mathbb{R})$ invariant set of Lemma 6.7. By EMM, the complement of this set is a finite union of proper affine invariant submanifolds.

The span of the (S_1, S_2) -recognizable twists gives a flat subbundle over \mathcal{M} minus this finite union of proper orbit closures, and this flat subbundle is not contained in $\ker(p) \cap T(\mathcal{M})$, because standard twists are never in $\ker(p)$.

By [Wri14, Theorem 5.1], any such flat subbundle must be all of $T(\mathcal{M})$. \square

6.3. (S_1, S_2) almost determines \mathcal{M} . The following result is not used in this paper. It is included to develop intuition, and because it may have applications in the future.

Theorem 6.9. *Fix finite sets $S_1, S_2 \subset \mathbb{R}$. In any fixed stratum there are at most finitely many affine invariant submanifolds \mathcal{M} such that (S_1, S_2) is \mathcal{M} -adapted.*

Proof. Otherwise, by EMM we would get an infinite sequence of \mathcal{M}_n for which (S_1, S_2) is \mathcal{M} -adapted and that are dense in some larger \mathcal{M} . By Lemma 6.7, there is a finite union of affine invariant submanifolds \mathcal{N}_i in \mathcal{M} , such that on $\mathcal{M} \setminus \cup \mathcal{N}_i$ the span of the (S_1, S_2) -recognizable twists is locally constant.

At most points of \mathcal{M}_n , the (S_1, S_2) -recognizable twists span $T(\mathcal{M}_n)$. This gives a contradiction, since the \mathcal{M}_n are dense in \mathcal{M} , and two different affine invariant submanifolds cannot have the same tangent space. (More precisely, given two nearby points $(X, \omega) \in \mathcal{M}_n$ and $(X', \omega') \in \mathcal{M}_{n'}$ the parallel translate from (X, ω) to (X', ω') of $T(\mathcal{M}_n)$ cannot be equal to $T(\mathcal{M}_{n'})$). \square

7. BOUNDARY (CONNECTED CASE)

In the first subsection we show that each point in $\partial\mathcal{M} \cap \overline{\mathcal{H}}_{conn}$ is contained in an affine invariant submanifold with the desired tangent space. In the second subsection, we bootstrap this to give a global statement about $\partial\mathcal{M} \cap \overline{\mathcal{H}}_{conn}$.

7.1. Affine invariant submanifolds containing limits. Although the rest of this section deals only with connected degenerations, the following proposition applies to all of $\overline{\mathcal{H}}$, not just $\overline{\mathcal{H}}_{conn}$.

Proposition 7.1. *Suppose that u is a (S_1, S_2) -recognizable twist at $(X, \omega) \in \overline{\mathcal{H}}$, and that $(X_n, \omega_n) \in \mathcal{H}$ is a sequence converging to (X, ω) . Then, for n large enough, u is in the span of the (S_1, S_2) -recognizable twists at (X_n, ω_n) .*

Let $V \subset H_1(X, \Sigma, \mathbb{Z})$ denote the space of vanishing cycles. Recall that the tangent space to the boundary stratum is naturally identified with $\text{Ann}(V) \subset H^1(X, \Sigma, \mathbb{C})$. It is this identification that we are using to consider u as a relative homology class at (X_n, ω_n) . In particular, note that at (X_n, ω_n) , by definition u is in $\text{Ann}(V)$.

Proof. This follows exactly as in the proof of Lemma 6.6, using Lemma 2.15. \square

Proposition 7.2. *Suppose that $(X_n, \omega_n) \in \mathcal{M}$ converge to $(X, \omega) \in \partial\mathcal{M} \cap \overline{\mathcal{H}}_{conn}$.*

Then (X, ω) is contained in an affine invariant submanifold whose tangent space at (X, ω) is given, for infinitely many n , by $\text{Ann}(V) \cap$

$T_{(X_n, \omega_n)}(\mathcal{M}) \subset H^1(X_n, \Sigma_n, \mathbb{C})$ via the isomorphism between the tangent space to the boundary stratum at (X, ω) and $\text{Ann}(V) \subset H^1(X_n, \Sigma_n, \mathbb{C})$.

Remark 7.3. It could be that (X, ω) is also contained in a smaller affine invariant submanifold even if all (X_n, ω_n) have orbit closure \mathcal{M} , although this is not the generic situation.

Proof. By Proposition 2.6, there is a neighborhood of 0 in $\text{Ann}(V)$ such that for ξ in this neighborhood and n large enough, $(X_n, \omega_n) + \xi$ and $(X, \omega) + \xi$ are well defined, and such that if $\xi_n \in \text{Ann}(V)$ are in this neighborhood and $\xi_n \rightarrow \xi$, then $(X_n, \omega_n) + \xi_n$ converges to $(X, \omega) + \xi$.

By Proposition 2.5, each subspace $T_{(X_n, \omega_n)}(\mathcal{M}) \cap \text{Ann}(V) \subset H^1(X_n, \Sigma_n, \mathbb{C})$ can be thought of as contained in a fixed vector space isomorphic to the tangent space of the boundary stratum. It suffices to pass to a subsequence along which these subspaces converge in this fixed vector space. Assume we have already passed to this subsequence.

Let U denote the set of all limits of all such sequences $(X_n, \omega_n) + \xi_n$, where $\xi_n \in \text{Ann}(V) \cap T_{(X_n, \omega_n)}(\mathcal{M})$ is in the neighborhood of 0 referred to above. Note that U is equal to (X, ω) plus all sufficiently small vectors in the subspace of $\text{Ann}(V)$ to which $T_{(X_n, \omega_n)}(\mathcal{M}) \cap \text{Ann}(V) \subset H^1(X_n, \Sigma_n, \mathbb{C})$ converges.

Let \mathcal{M}' be the smallest closed $GL(2, \mathbb{R})$ invariant set containing U . Because all closed $GL(2, \mathbb{R})$ invariant sets are finite unions of affine invariant submanifolds, and because U is linear, we see that \mathcal{M}' is a single affine invariant submanifold.

Let (S_1, S_2) be both \mathcal{M} -adapted and \mathcal{M}' -adapted. By Proposition 6.8, away from a finite union of proper affine invariant submanifolds, the tangent space to \mathcal{M}' is spanned by (S_1, S_2) -recognizable twists.

Pick a specific sequence ξ_n so that the limit of the $(X_n, \omega_n) + \xi_n$ is a surface (X', ω') not contained in this finite union of proper affine invariant submanifolds.

The tangent space to \mathcal{M}' at (X', ω') is spanned by (S_1, S_2) -recognizable twists. By Proposition 7.1, all the (S_1, S_2) -recognizable twists at (X', ω') are recognizable at $(X_n, \omega_n) + \xi_n$ as well. Hence they are all in $\text{Ann}(V) \cap T_{(X_n, \omega_n)}(\mathcal{M})$.

We have shown that the tangent space to \mathcal{M}' at (X', ω') is contained in $\text{Ann}(V) \cap T_{(X_n, \omega_n)}(\mathcal{M})$ for all n large enough. Let us now address the opposite containment. Suppose $T(\mathcal{M}')$ was strictly contained in $\text{Ann}(V) \cap T_{(X_n, \omega_n)}(\mathcal{M})$ for all n . In this case, pick $\xi'_n \in \text{Ann}(V) \cap T_{(X_n, \omega_n)}(\mathcal{M})$ so that ξ'_n converges to $\xi \notin T(\mathcal{M}')$. The definition of \mathcal{M}' gives that for t small enough, $(X', \omega') + t\xi \in \mathcal{M}$. Hence $\xi \in T(\mathcal{M}')$, a contradiction. \square

Remark 7.4. A variant of the above approach allows the use of sets (S_1, S_2) that are \mathcal{M} -adapted without knowing whether they are \mathcal{M}' -adapted. The proof of Proposition 6.8 shows that for any affine invariant submanifold \mathcal{M}' , and arbitrary (S_1, S_2) , away from a finite union of proper affine invariant submanifolds the span of the (S_1, S_2) -recognizable twists contains $T(\mathcal{M}')$. In the above proof, one can then pick (X', ω') such that the span of the (S_1, S_2) -recognizable twists contains $T(\mathcal{M}')$. This strategy will be required in Section 8.

7.2. Global boundary. We now rule out the possibility that infinitely many boundary orbit closures \mathcal{M}' occur in Proposition 7.2. This will conclude the proof of the first part of Theorem 2.9, which appeared in the introduction as Theorem 1.1.

Theorem 7.5. *Only finitely many affine invariant submanifolds arise in Proposition 7.2.*

That is, there is a finite list of affine invariant submanifolds in the boundary of \mathcal{M} such that if $(X_n, \omega_n) \in \mathcal{M}$ converge to $(X, \omega) \in \partial\mathcal{M} \cap \overline{\mathcal{H}_{conn}}$, then (X, ω) is contained in one of these finitely many affine invariant submanifolds, and the tangent space to this affine invariant submanifold can be computed as in Proposition 7.2.

Proof. Suppose in order to find a contradiction that infinitely many affine invariant submanifolds arise in Proposition 7.2. Pick a sequence \mathcal{M}_k of them that all have maximal dimension and all lie in the same boundary stratum.

By EMM, there are $(X^{(k)}, \omega^{(k)}) \in \mathcal{M}_k'$ that converge to some $(X^{(\infty)}, \omega^{(\infty)})$ that has larger dimensional orbit closure \mathcal{M}'_∞ .

Pick a sequence $(X_n^{(k)}, \omega_n^{(k)}) \rightarrow (X^{(k)}, \omega^{(k)})$ with $(X_n^{(k)}, \omega_n^{(k)}) \in \mathcal{M}$ and such that $T(\mathcal{M}_k') = \text{Ann}(V) \cap T_{(X_n^{(k)}, \omega_n^{(k)})}(\mathcal{M})$.

Some “diagonal” sequence $(X_{n_k}^{(k)}, \omega_{n_k}^{(k)})$ then converges to $(X^{(\infty)}, \omega^{(\infty)})$. Thus, for infinitely many n , the tangent space to both \mathcal{M}_k' and \mathcal{M}'_∞ can be identified with

$$\text{Ann}(V) \cap T_{(X_{n_k}^{(k)}, \omega_{n_k}^{(k)})}(\mathcal{M}),$$

which is a contradiction because \mathcal{M}_k' and \mathcal{M}'_∞ have different dimensions.

Proposition 2.5 allows us to consider V without worrying about dependence on k or n_k . \square

Proof of Theorem 2.9, first part. Proposition 7.2 shows that each point in the boundary appears in an affine invariant submanifold with the

desired tangent space. The proof additionally shows that the set of limit points are dense in the affine invariant submanifold.

Theorem 7.5 shows that only finitely many affine invariant submanifolds appear in this way. Hence these finitely many affine invariant submanifolds constitute the boundary of \mathcal{M} in the given boundary stratum. \square

8. MULTICOMPONENT CASE

In the first subsection, we prove the second statement in Theorem 2.9, which is the direct generalization of Proposition 7.2 to multicomponent surfaces, by explaining what modifications to the proof of Proposition 7.2 are required. In the next two subsections, we assume multicomponent EMM is true, and prove Theorem 2.12. As in the connected case, this involves first showing that all limits are contained in an affine invariant submanifold with the expected tangent space (as in Proposition 7.2) and then showing that only finitely many such affine invariant submanifolds occur (as in Theorem 7.5).

8.1. Projection onto connected components. Consider a multicomponent limit (X', ω') of a sequence $(X_n, \omega_n) \in \mathcal{M}$. Say π is the projection onto a connected component. (So π just forgets the other connected components.)

As in the proof of Proposition 7.2, after passing to an appropriate subsequence, we consider the smallest affine invariant submanifold \mathcal{M}' that contains all $\pi(X', \omega')$, where (X', ω') is a limit of $(X_n, \omega_n) + \xi_n$, as in the proof of Proposition 7.2.

As in that proposition, we can assume (X', ω') has the property that $T(\mathcal{M}')$ is spanned by (S_1, S_2) -recognizable twists at $\pi(X', \omega')$. We now claim that for every (S_1, S_2) -recognizable twist u on $\pi(X', \omega')$, there is a (S_1, S_2) -recognizable twist on (X', ω') which restricts to u .

This is a consequence of simple linear algebra. Let \mathcal{E} be a set of homogeneous linear equations in variables t_1, \dots, t_n . Suppose that t_1, \dots, t_k satisfy all the linear equations in \mathcal{E} that only involve the first k variables. Then there are values of t_{k+1}, \dots, t_n so that t_1, \dots, t_n satisfy all linear equations in \mathcal{E} .

Thus, also as in Proposition 7.2, we see that $T(\mathcal{M}')$ is contained in $\pi(\text{Ann}(V) \cap T(\mathcal{M}))$. The proof that $T(\mathcal{M}')$ cannot be smaller than $\pi(\text{Ann}(V) \cap T(\mathcal{M}))$ is the same as in the connected case (Proposition 7.2).

8.2. Results using multicomponent EMM: Analogue of Proposition 7.2. Note that every affine invariant submanifold is a product of irreducible affine invariant submanifolds.

We follow the argument in the connected case, however some complications arise. As in the proof of Proposition 7.2, we consider the set U of limits of sequences $(X_n, \omega_n) + \xi_n$, where $\xi_n \in \text{Ann}(V) \cap T_{(X_n, \omega_n)}(\mathcal{M})$ are small enough. We may assume that U is a small locally linear set in the boundary. The main complication is the possibility, which is ruled out only at the conclusion of our arguments, that the smallest closed $GL(2, \mathbb{R})$ invariant set containing U is not an affine invariant submanifold.

To be concrete, let us mention that one such bad situation would be if the smallest closed invariant set containing U was the set of all surfaces in $\mathcal{H}(2) \times \mathcal{H}(2)$ where the two components have equal area. The following is an example of a locally linear set U with exactly this bad property.

Example 8.1. Choose $(X, \omega) \in \mathcal{H}(2)$ with dense orbit, and consider $v \in H^1(X, \mathbb{R})$ that is symplectically orthogonal to $\text{Re}(\omega)$ and $\text{Im}(\omega)$. Note that for any small such v , $(X, \omega) + v$ and (X, ω) have the same area.

Define U to be the subset of $\mathcal{H}(2) \times \mathcal{H}(2)$ given by $((X, \omega), (X, \omega)) + \xi$, where ξ ranges over a neighborhood of zero in

$$\text{span}_{\mathbb{C}}((\text{Re}(\omega), \text{Re}(\omega)), (\text{Im}(\omega), \text{Im}(\omega)), (0, v)).$$

Note that every element of U has that the two components have equal area. Assuming multicomponent EMM, one can show that the smallest $GL(2, \mathbb{R})$ invariant set containing U is the set of all surfaces in $\mathcal{H}(2) \times \mathcal{H}(2)$ where the two components have equal area. (Indeed, by Conjecture 2.10, the smallest closed invariant set containing U must be the subset of some affine invariant submanifold where the two components have equal area. By construction and part (2) of the conjecture, this affine invariant submanifold can't be irreducible. Since by construction its projection to each coordinate must be all of $\mathcal{H}(2)$, we see that it must be $\mathcal{H}(2) \times \mathcal{H}(2)$.)

Our first lemma shows that irreducible affine invariant submanifolds do not contain any such strange sets.

Lemma 8.2. *Let \mathcal{M} be an irreducible affine invariant submanifold. Then any closed $GL(2, \mathbb{R})$ invariant subset of \mathcal{M} is a finite union of affine invariant submanifolds.*

Proof. In an irreducible affine invariant submanifold, the ratio of areas of components is constant. By scaling factors, it can be assumed that

the constants are one, and then the statement is given by multicomponent EMM. \square

Lemma 8.3. *Let \mathcal{M} be an irreducible affine invariant submanifold. Any flat subbundle of $T(\mathcal{M})$ not contained in $\ker(p)$ is equal to all of $T(\mathcal{M})$.*

Proof. This follows as in [Wri14, Theorem 5.1]. Irreducibility is only used so that the geodesic flow is hyperbolic. \square

Lemma 8.4. *Let \mathcal{M} be an irreducible affine invariant submanifold. Let (S_1, S_2) be arbitrary finite sets. Then there is a dense open $GL(2, \mathbb{R})$ invariant subset of \mathcal{M} on which the span of the (S_1, S_2) -recognizable twists contains $T(\mathcal{M})$.*

Proof. Given the previous lemma, this follows from the proof of Proposition 6.8, following Remark 7.4. \square

We now pause to outline the remainder of this subsection. Keeping in mind Example 8.1, we cannot assume at this point that the smallest closed $GL(2, \mathbb{R})$ set containing U is an affine invariant submanifold rather than a nonlinear object. We circumvent this issue by first working with the span of recognizable twists. Lemma 8.5 gives a condition that can be used to show this span is at least as large as the tangent space to the smallest affine invariant submanifold containing U , and in Lemma 8.6 this condition is established almost everywhere on U . Finally, in the course of the proof of Proposition 8.7, it is shown that U is an open subset of the span of the recognizable twists, which will allow us to conclude our analysis.

Lemma 8.5. *Let $\mathcal{M}_1, \dots, \mathcal{M}_k$ be irreducible affine invariant submanifolds. Set $\mathcal{M} = \mathcal{M}_1 \times \dots \times \mathcal{M}_k$.*

Let $(X, \omega) \in \mathcal{M}$, and assume that there is an (X', ω') in the orbit closure of (X, ω) with the following properties.

First, no cylinder in any factor of (X', ω') has a parallel cylinder on any other factor with ratio of circumferences in S_1 . Second, $\pi_i(X', \omega')$ has dense $GL(2, \mathbb{R})$ orbit in \mathcal{M}_i for each i .

Then the span of the (S_1, S_2) -recognizable twists at (X, ω) contains $T(\mathcal{M})$.

Here $\pi_j : \mathcal{M} \rightarrow \mathcal{M}_j$ is the projection into the j -th factor. By a cylinder on the j -th factor of (X, ω) , we mean a cylinder on $\pi_j(X, \omega)$.

Proof. By Lemma 8.4, the span of the (S_1, S_2) -recognizable twists contains $T(\mathcal{M})$ at (X', ω') . Since the orbit of (X, ω) accumulates at (X', ω') , by Lemma 6.6 we have that the same holds for (X, ω) . \square

Lemma 8.6. *Let $\mathcal{M}_1, \dots, \mathcal{M}_k$ be irreducible affine invariant submanifolds. Set $\mathcal{M} = \mathcal{M}_1 \times \dots \times \mathcal{M}_k$.*

Suppose U is a subset of \mathcal{M} that is equal to an open subset of a vector subspace in period coordinates, and suppose that \mathcal{M} is the smallest affine invariant submanifold containing U . Let $S_1 \subset \mathbb{R}$ be any finite set. Then for almost every surface (X, ω) in U , almost every point (X', ω') in the orbit closure of (X, ω) has the following properties.

First, no cylinder in any factor of (X', ω') has a parallel cylinder on any other factor with ratio of circumferences in S_1 . Second, $\pi_i(X', \omega')$ has dense $GL(2, \mathbb{R})$ orbit in \mathcal{M}_i for each i .

Proof. The case $k = 1$ follows from Lemma 8.2. Assume $k = 2$.

Let U_1 be the subset of (X, ω) in U such that the orbit closure $\pi_i(X, \omega)$ is \mathcal{M}_i for $i = 1, 2$. This set U_1 has full measure, because it is the complement of a countable number of linear subspaces (one for each affine invariant submanifolds of each \mathcal{M}_i).

Case 1: For almost every $(X, \omega) \in U$, the smallest affine invariant submanifold containing (X, ω) is reducible. The result follows, because for any such $(X, \omega) \in U_1$, the orbit closure of (X, ω) is the subset of $\mathcal{M}_1 \times \mathcal{M}_2$ where the two components have the same ratio of area as (X, ω) .

Case 2: A positive measure subset of $(X, \omega) \in U$ is contained in an irreducible affine invariant submanifold. Let U_2 be the subset of U_1 with this property. For $(X, \omega) \in \mathcal{M}_1 \times \mathcal{M}_2$, let $r(X, \omega)$ be the result of scaling the two factors so that they both have equal area. By the countability statement in multicomponent EMM, there is an affine invariant submanifold $\mathcal{N} \subset \mathcal{M}_1 \times \mathcal{M}_2$ such that the following holds. For a positive measure set $U_3 \subset U_2$, for every $(X, \omega) \in U_3$ the orbit closure of $r(X, \omega)$ is the subset of \mathcal{N} where the factors both have equal area. Because $r^{-1}(\mathcal{N})$ is an analytic submanifold, since it contains a positive measure subset of U , it must in fact contain all of U .

Note that the ratio of the area of the two components is a non-constant analytic function on U . (If it were constant, then U would be contained in a scaling of \mathcal{N} , showing that $\mathcal{M}_1 \times \mathcal{M}_2$ is not the smallest affine invariant submanifold containing U .)

Let \mathcal{N}^{gen} be the set surfaces in \mathcal{N} on which all parallel cylinders are in fact \mathcal{N} -parallel. Note that \mathcal{N}^{gen} has full measure. Let S'_1 be the set of ratios of circumferences of \mathcal{N} -parallel cylinders on surfaces in \mathcal{N} . This set is countable. (In fact, a generalized version of the Cylinder Finiteness Theorem applies here and gives that S'_1 is finite, but we only need countability here and that is easy.)

Let (X, ω) be any surface in U_3 where the square root of the ratio of the areas of the two factors is not an element of S_1 divided by an element of S'_1 . By the definition of U_3 , the orbit closure of (X, ω) is \mathcal{N} with the components rescaled, and any surface (X', ω') in the $GL(2, \mathbb{R})$ orbit closure of (X, ω) with $r(X', \omega') \in \mathcal{N}^{gen}$ has the desired property that no cylinder on $\pi_1(X', \omega')$ is parallel to a cylinder in $\pi_2(X', \omega')$ with ratio of circumferences in S_1 .

Now, we will derive the result for $k > 2$ from the result for $k = 2$. For each i, j , let $U_{i,j}$ be the subset of $(X, \omega) \in U$ for which almost every surface (X', ω') in the orbit closure of (X, ω) has that $\pi_i(X', \omega')$ has dense $GL(2, \mathbb{R})$ orbit in \mathcal{M}_i , and $\pi_j(X', \omega')$ has dense $GL(2, \mathbb{R})$ orbit in \mathcal{M}_j , and that no cylinder on $\pi_i(X', \omega')$ is parallel to a cylinder in $\pi_j(X', \omega')$ with ratio of circumferences in S_1 . The lemma for $k = 2$ gives that $U_{i,j}$ has full measure. The intersection of the $U_{i,j}$ gives the full measure set desired. \square

Proposition 8.7. *Suppose that $(X_n, \omega_n) \in \mathcal{M}$ converge to $(X, \omega) \in \partial\mathcal{M}$.*

Then (X, ω) is contained in an affine invariant submanifold \mathcal{M}' whose tangent space at (X, ω) is given by, for infinitely many n , $\text{Ann}(V) \cap T_{(X_n, \omega_n)}(\mathcal{M}) \subset H^1(X_n, \Sigma_n, \mathbb{C})$ via the isomorphism between the tangent space to the boundary stratum at (X, ω) and $\text{Ann}(V) \subset H^1(X_n, \Sigma_n, \mathbb{C})$.

Also, any finite sets (S_1, S_2) that are \mathcal{M} -adapted are also \mathcal{M}' adapted.

Proof. Define U as in the proof of Proposition 7.2, and let \mathcal{M}' be the smallest affine invariant submanifold containing U .

Pick (S_1, S_2) that is \mathcal{M} -adapted.

By Lemmas 8.5 and 8.6, for some limit $(X', \omega') \in U$ the span of the (S_1, S_2) -recognizable twists contains $T(\mathcal{M}')$. By Proposition 7.1, all these twists are contained in the span of the (S_1, S_2) -recognizable twists at $(X_n, \omega_n) + \xi_n$ for n large enough, and hence we see that they are contained in $T(\mathcal{M}) \cap \text{Ann}(V)$. We get that $T(\mathcal{M}') \subset T(\mathcal{M}) \cap \text{Ann}(V)$.

The other containment proceeds just as in the proof of Proposition 7.2.

Proposition 7.1 also shows that (S_1, S_2) is \mathcal{M}' -adapted, because any (S_1, S_2) -recognizable twist u at a surface in $\partial\mathcal{M}$ also gives a (S_1, S_2) -recognizable twist at a nearby surface in \mathcal{M} . Since (S_1, S_2) is \mathcal{M} -adapted, we have $u \in T(\mathcal{M})$. By definition, we have that $u \in \text{Ann}(V)$. Hence $u \in T(\mathcal{M}) \cap \text{Ann}(V) = T(\mathcal{M}')$. \square

8.3. Results using multicomponent EMM: Analogue of Theorem 7.5. Now we move on toward a version of Theorem 7.5. The new danger in the multicomponent case is that a sequence of affine

invariant submanifolds might not equidistribute to a larger affine invariant submanifold. For example, if \mathcal{M}_n is the locus in $\mathcal{H}(2) \times \mathcal{H}(2)$ where the second component is n times the first, then the \mathcal{M}_n diverge in $\mathcal{H}(2) \times \mathcal{H}(2)$.

Let $\mathcal{H}_i, i = 1, \dots, k$ be strata of connected translation surfaces, and say \mathcal{M} is an affine invariant submanifold of $\mathcal{H}_1 \times \dots \times \mathcal{H}_k$. Let $A(\mathcal{M})$ be the set of $a \in \mathbb{R}$ such that there are $1 \leq i, j \leq k, i \neq j$ such that in \mathcal{M} , the area of the i -th component is always a times the area of the j -th component. So in the example above, $A(\mathcal{M}_n) = \{n, 1/n\}$.

Lemma 8.8. *For any finite sets (S_1, S_2) and $d > 0$, there is a finite set $A(S_1, S_2, d) \subset \mathbb{R}$ such that the following holds. Suppose \mathcal{M} is an affine invariant submanifold in a stratum of dimension at most d , and that (S_1, S_2) is \mathcal{M} -adapted. Then $A(\mathcal{M}) \subset A(S_1, S_2, d)$.*

Proof. For simplicity, suppose \mathcal{M} consists of surfaces with two components. Pick any surface $(X, \omega) \in \mathcal{M}$, and consider the set $\mathcal{C} = \{C_1, \dots, C_k\}$ of cylinders in some direction. Note $k < d$. (This is the only way d will be used.)

Consider the equivalence relation on \mathcal{C} generated by C_i equivalent to C_j if the ratio of circumferences is in S_1 . Let $\mathcal{C}' \subset \mathcal{C}$ be an equivalence class for this equivalence relation. Note that the ratio of circumferences of any two cylinders in \mathcal{C}' is in a finite set depending only on S_1 and d . Note also that \mathcal{C}' supports at least one (S_1, S_2) -recognizable twist, corresponding to $t_i = 0$ for cylinders in $\mathcal{C} \setminus \mathcal{C}'$, and t_i equal to the modulus for the cylinders in \mathcal{C}' .

Let α_i be the core curve of C_i . Without loss of generality, assume $\mathcal{C}' = \{C_1, \dots, C_{k'}\}$.

By definition, the (S_1, S_2) -recognizable twists in \mathcal{C}' are the twists of the form

$$u = \sum_{i=1}^{k'} t_i c_i I_\Sigma(\alpha_i),$$

where the t_i satisfy all linear equations with coefficients in S_2 that the moduli m_i satisfy. There are only finitely many systems of linear equations in S_2 with at most d variables. For each one of them that admits a solution $(t_1, \dots, t_{k'})$ with all $t_i > 0$, pick a solution $t_i > 0$.

Now, we use the fixed solution $(t_1, \dots, t_{k'})$ to the set of linear equations defining the (S_1, S_2) -recognizable twists in \mathcal{C} . Say that C_1, \dots, C_ℓ are contained in the first component, and $C_{\ell+1}, \dots, C_{k'}$ in the second. Then, applying the cylinder stretch arising from u (an appropriate complex multiple of u that changes area), we see that the area of the first

component has derivative $\sum_{i=1}^{\ell} t_i c_i$, and the area of the second component has derivative $\sum_{i=\ell+1}^{k'} t_i c_i$. Since the ratios of areas is constant, this ratio must be $\sum_{i=1}^{\ell} t_i c_i$ divided by $\sum_{i=\ell+1}^{k'} t_i c_i$. \square

Lemma 8.9. *Let A be a finite set, and let \mathcal{H} be a stratum of multicomponent translation surfaces. Suppose that $\mathcal{M}_i, i = 1, 2, \dots$, is an infinite sequence of distinct affine invariant submanifolds of fixed dimension d of \mathcal{H} with $A(\mathcal{M}_i) \subset A$.*

Then the closure of the union of the \mathcal{M}_i contains an affine invariant submanifold of dimension larger than d .

Proof. Passing to a subsequence we may assume that all \mathcal{M}_i have the same pairs of components with the same fixed area ratios. In other words, \mathcal{H} is the product of strata of (possibly multicomponent) surfaces $\mathcal{H} = \mathcal{H}_1 \times \dots \times \mathcal{H}_p$, and each \mathcal{M}_i is a product of irreducibles, $\mathcal{M}_i = \mathcal{M}_{i,1} \times \dots \times \mathcal{M}_{i,p}$. Here $\mathcal{M}_{i,j} \subset \mathcal{H}_j$, and the ratios of the components in each $\mathcal{M}_{i,j}$ does not depend on j .

For each j , since all $\mathcal{M}_{i,j}$ live in the subset of \mathcal{H}_j where the ratios of areas are fixed independent of i , the closure of $\cup_j \mathcal{M}_{i,j}$ is a finite union of affine invariant submanifolds. Let \mathcal{N}_j be the component of largest dimension.

The union of the closure of the \mathcal{M}_i contains $\mathcal{N}_1 \times \dots \times \mathcal{N}_p$, and it is easy to see that this has dimension larger than d . \square

Proof of Theorem 2.12. Given Proposition 8.7, to conclude the proof of Theorem 2.12 it suffices to prove a version of Theorem 7.5 in the multicomponent situation.

So we assume to the contrary, that for some fixed boundary stratum \mathcal{H}' , there are infinitely many affine invariant submanifolds \mathcal{M}_i arising as in Proposition 8.7. Pass to a subsequence where the \mathcal{M}_i have maximal dimension d . The previous two lemmas give that the closure of the union of the \mathcal{M}_i contains an affine invariant submanifold of dimension larger than d , and so we can conclude the proof exactly as in the proof of Theorem 7.5. \square

9. BASIC BOUNDARY THEORY FOR STRATA

In the first subsection we prove Propositions 2.4 and 2.5. In the second we prove Proposition 2.6. In the third, we show that the notion of convergence defined in Section 2 can be rephrased in terms of the Deligne-Mumford compactification of $\mathcal{M}_{g,s}$.

9.1. Proof of Proposition 2.5. The next lemma considers a stratum \mathcal{H} in the boundary of some other stratum \mathcal{H}' .

Lemma 9.1. *Let \mathcal{H}' and \mathcal{H} be two strata of multicomponent translation surfaces, and let $(X, \omega, \Sigma) \in \mathcal{H}$. Let U be small neighborhood of Σ , such that each component of U is homeomorphic to a disk. Fix $\delta > 0$ and a neighborhood \mathcal{O} of ω in the space of differential 1-forms on $X \setminus U$. Let \mathcal{U} be the neighborhood in \mathcal{H}' of (X, ω, Σ) that consists of $(X', \omega', \Sigma') \in \mathcal{H}'$ such that there is a map $g : X \setminus U \rightarrow X'$ that is a diffeomorphism onto its range, such that*

- (1) $g^*(\omega') \in \mathcal{O}$ and
- (2) the injectivity radius of points not in the image of g is at most δ .

Let Δ be the length of the shortest saddle connection on (X, ω, Σ) . Then for U, δ, \mathcal{O} small enough, any saddle connection on (X', ω', Σ') of length less than $\Delta/2$ is contained in the δ -thin part (the subset with injectivity radius at most δ).

Proof. We may pick U to be a union of small balls about points of Σ small enough that any path between different components of U on (X, ω, Σ) has length at least $\frac{3}{4}\Delta$. If \mathcal{O} is chosen small enough, then any path between different components of $X' \setminus g(X \setminus U)$ has length at least $\Delta/2$, which gives the result since $X' \setminus g(X \setminus U)$ is contained in the δ -thin part by supposition. \square

Proof of Propositions 2.4 and 2.5. We may assume that the U_n are homeomorphic to disks and the $U_{n-1} \setminus U_n$ are homeomorphic to annuli.

Given $g_n : X \setminus U_n \rightarrow X_n$, define the collapse map f_n to be the inverse of g_n on $X \setminus g_n(U_{n-1})$, and to map $X_n \setminus g_n(X \setminus U_n)$ to the point $\Sigma \cap U_n$. The remainder of X_n is a collection of annuli on which f_n can be defined to interpolate.

By Lemma 9.1, for all $\delta > 0$ small enough and n large enough,

$$V_n = \ker(f_n : H_1(X_n, \Sigma_n, \mathbb{Z}) \rightarrow H_1(X, \Sigma, \mathbb{Z}))$$

is spanned by the relative homology of the δ -thin part. (Indeed, the set $X_n \setminus g_n(X \setminus U_n)$ that is collapsed can be assumed to be in the thin part, and Lemma 9.1 says that the thin part isn't larger.) Also for n large enough, V_n is locally constant: there is a neighborhood \mathcal{U} of (X, ω, Σ) given by Lemma 9.1 that contains all $(X_n, \omega_n, \Sigma_n)$ for n sufficiently large, on which the homology of the δ' -thin part is invariant under parallel transport. \square

9.2. Proof of Proposition 2.6: How strata fit together. Suppose that $(X_n, \omega_n, \Sigma_n) \in \mathcal{H}'$ converges to $(X, \omega, \Sigma) \in \mathcal{H}$, and let $V \subset H_1(X_n, \Sigma_n, \mathbb{Z})$ be the space of vanishing cycles.

Lemma 9.2. *$\text{Ann}(V)$ is naturally identified with the tangent space of the boundary stratum \mathcal{H} via the collapse maps f_n .*

Proof. By the definition of V as the kernel of the action of the collapse map on relative homology, and because relative cohomology is the dual to relative homology, the induced map f_n^* on relative cohomology is an isomorphism from $H^1(X, \Sigma, \mathbb{C})$ to $\text{Ann}(V)$. \square

The remainder of Proposition 2.6 is given by the following.

Proposition 9.3. *There exists a neighborhood of 0 in $\text{Ann}(V)$ and an $n_0 > 0$ such that for ξ, ξ_n in this neighborhood with $\xi_n \rightarrow \xi$, the paths $(X_n, \omega_n, \Sigma_n) + t\xi_n$ and $(X, \omega, \Sigma) + t\xi$ are well defined for $t \in [0, 1]$ and $n > n_0$, and $(X_n, \omega_n, \Sigma_n) + \xi_n \rightarrow (X, \omega, \Sigma) + \xi$.*

It is plausible that part of this (that the paths are well defined) could be derived using the AGY norm [AGY06], in particular using [AG13, Proposition 5.5]. Here we will instead sketch a proof using Čech cohomology.

The use of Čech cohomology to describe the tangent space to Teichmüller space is standard and is sometimes called Kodaira-Spencer deformation theory. See, for example, [IT92, Section 7.2.4] for an introduction. Čech cohomology has also been previously used to describe deformations of translation surfaces. See, for example, [HM79, Section 2] and [McM13].

Given a finite open cover $\mathcal{U} = \{U_1, \dots, U_p\}$ of a translation surface (X, ω, Σ) , we consider 1-cocycles w of \mathcal{U} with coefficients in the sheaf of locally constant functions. Such a w consists of a complex number $w_{i,j} \in \mathbb{C}$ for each $1 \leq i, j \leq p$ with $U_i \cap U_j \neq \emptyset$. These numbers are required to satisfy $w_{i,j} = -w_{j,i}$ for all i, j , and $w_{i,k} = w_{i,j} + w_{j,k}$ for all i, j, k with $U_i \cap U_j \cap U_k \neq \emptyset$. We define $\|w\| = \max |w_{i,j}|$.

We consider deformations of (X, ω, Σ) defined by sliding each piece U_j over the piece U_i in direction $w_{i,j}$. See Figure 9.1. Such a deformation is not always well defined. For example, as the pieces slide more and more over each other, a zero of ω can hit the boundary of some $U_{i,j}$, preventing that $U_{i,j}$ from being slid any further (as in Figure 9.2, left); or, if pieces are being slid away from each other, a “hole” in the surface may appear (as in Figure 9.2, right).

The following definition and lemma provide a precise definition of such deformations and a proof that they are well defined when w is small enough and the open cover is nice enough.

Definition 9.4. Define a finite open cover $\mathcal{U} = \{U_1, \dots, U_p\}$ of a translation surface (X, ω, Σ) to be decent if each point of Σ is contained in

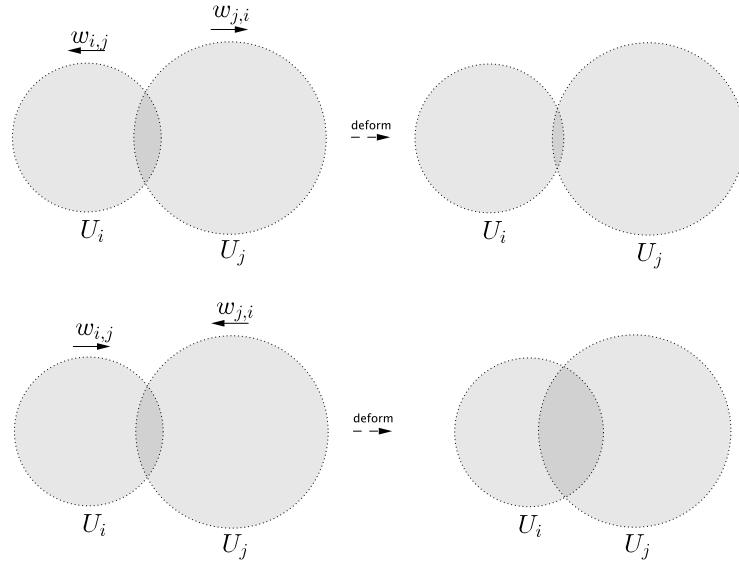


FIGURE 9.1

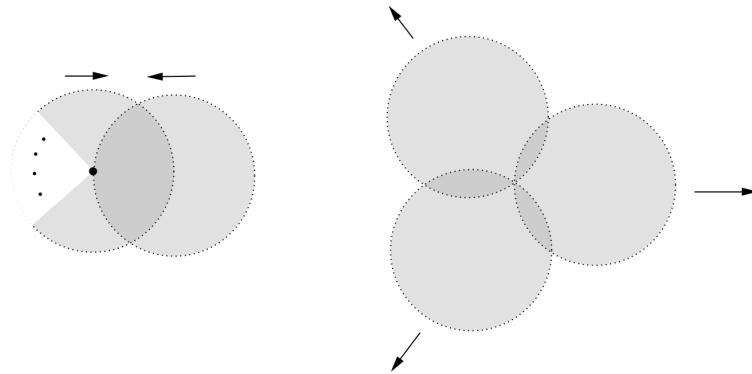


FIGURE 9.2

exactly one U_i and is disjoint from the closures of every other U_j , and the closures of U_i and U_j are disjoint whenever U_i and U_j are disjoint. Define \mathcal{U} to be ε decent if the following conditions hold.

- (1) If U_i and U_j are disjoint, the distance between them is greater than 2ε .
- (2) Every point of Σ is contained in a unique U_i , and its distance to every other U_j is greater than 2ε .
- (3) For every point $x \in X$, there is some U_i that contains all points of distance at most ε from U_i .

Note that every decent cover is ε decent for some $\varepsilon > 0$, and that the definition allows a U_i to contain more than one point of Σ .

If x is a point on (X, ω, Σ) not in Σ , and $e \in \mathbb{C}$, then $x + e$ is defined to be the result of straight line flow in direction $e/|e|$ for time $|e|$; if the trajectory of straight line flow hits a point of Σ before time $|e|$, then $x + e$ is undefined.

Lemma 9.5. *Let $\mathcal{U} = \{U_1, \dots, U_p\}$ be an ε decent open cover of (X, ω, Σ) . Let w be a 1-cocycle of \mathcal{U} with coefficients in the sheaf of locally constant functions, and assume $\|w\| < \varepsilon$.*

Consider the disjoint union of the U_i with the following identifications. A point $x_i \in U_i$ is identified to $x_j \in U_j$ if $U_i \cap U_j \neq \emptyset$ and $x_i + w_{i,j}$ is defined and equal to x_j . This defines a translation surface in the same stratum as (X, ω, Σ) .

Furthermore, the path of translation surfaces corresponding to $tw, t \in [0, 1]$, is linear in period coordinates.

Note that the lemma allows the U_i to be rather arbitrary: they may have complicated boundaries, and neither the U_i nor their intersections are assumed to be simply connected or even connected.

Proof. Say that $x_i \in U_i$ is identified to $x_j \in U_j$ if $U_i \cap U_j \neq \emptyset$ and $x_i + w_{i,j}$ is defined and equal to x_j . (By definition, this implies there is a straight line segment from x_i to x_j on (X, ω, Σ) with holonomy $w_{i,j}$. This straight line segment is not required to stay in $U_i \cup U_j$.) A priori, there is the possibility that $x_i \in U_i$ and $x_j \in U_j$ could have the same image in the quotient by these identification without being identified to each other. This will ruled out shortly when we show the relationship “is identified to” is transitive.

Suppose $x \in U_i$ has distance less than $\|w\|$ to a point of Σ . By Definition 9.4(2) this point of Σ lies in U_i , and the distance between x and any $U_j, j \neq i$ is greater than $\|w\|$. Hence the $\|w\|$ neighborhood of every point of Σ is contained in a single U_i , and no point of this neighborhood is identified to any point of any $U_j, j \neq i$. Note also that for any point x of distance at least $\|w\|$ from Σ , $x + e$ is defined whenever $|e| \leq \|w\|$.

Suppose $x_i \in U_i$ is identified to $x_j = x_i + w_{i,j} \in U_j$, and x_j is identified to $x_k = x_j + w_{j,k} \in U_k$. Then the distance from U_i to U_k is at most $|w_{i,j}| + |w_{j,k}| \leq 2\|w\|$, so $U_i \cap U_k \neq \emptyset$ by Definition 9.4(1). Consider the flat geodesic γ from x_i to x_j homotopic rel endpoints to the concatenation of the straight line flow path from x_i to x_j followed by the straight line flow path from x_j to x_k . By the previous comment, both x_i and x_k have distance greater than $\|w\|$ from Σ . Since γ has

length at most $2\|w\|$, this gives that γ is a straight line segment from x_i to x_k . This shows that $x_k = x_i + w_{i,k}$, so we conclude that x_i is identified to x_k .

Gluing pieces of translation surfaces via translations will always yield a translation surface, so it remains only to check that this translation surface is closed (unlike Figure 9.2, right). By Definition 9.4(3), for each point x in the boundary of a U_i , there is some U_j that contains all points on (X, ω, Σ) of distance at most $\|w\|$ from x . After the identifications described, this U_j will be glued to U_i in such a way to cover the point x on the boundary of U_i . This shows the identification gives a closed surface.

To verify that the path of translation surfaces obtained is linear, it suffices to show that if $U_i \cap U_j \neq \emptyset$, then for any path joining a point of $\Sigma \cap U_i$ with a point of $\Sigma \cap U_j$ the period of this path varies linearly. This is the case by definition. \square

Lemma 9.6. *For each translation surface (X, ω, Σ) , there is a decent open cover such that the derivatives of deformations of (X, ω, Σ) that arise as above from 1-cocycles on this cover span the tangent space to the stratum at (X, ω, Σ) . This open cover may be chosen so that all points of Σ are contained in a single element of the open cover.*

The proof follows from general theory, but we give a sketch of a concrete approach.

Proof. We begin with the first statement. Pick a maximal spanning subtree of a triangulation of the surface. A cover given by neighborhoods of half edges, and the complement of a neighborhood of the spanning tree, is easily seen to work. The periods of the edges in the spanning tree can be changed independently by deformations given by cocycles.

To arrange for all elements of Σ to be contained in a single element of the open cover, we may modify the cover $\mathcal{U} = \{U_1, \dots, U_p\}$ as follows. We use the fact that each element of Σ is contained in a single U_i and is disjoint from the $\overline{U_j}$, $j \neq i$. First, add an open set U'_0 that consists of points of distance less than $2r$ from Σ , where $r > 0$ is very small. Then remove from each U_i the set of distance at most r from Σ to get a new open set U'_i . Then $\{U'_0, U'_1, \dots, U'_p\}$ will be the new open cover.

For any cocycle w supported on the original open cover, we can find a new cocycle w' supported on the new open cover by $w'_{i,j} = w_{i,j}$ if $i, j \neq 0$ and $w'_{0,i} = w'_{i,0} = 0$. The cocycles w and w' describe exactly the same deformation. \square

Lemma 9.7. *Fix an ε decent cover of (X, ω, Σ) , such that the derivatives of deformations of (X, ω, Σ) that arise as above from 1-cocycles on this cover span the tangent space to the stratum at (X, ω, Σ) , and so that Σ is contained in a single element of the cover. For n sufficiently large, the pull back of this cover to $(X_n, \omega_n, \Sigma_n)$ is $\varepsilon/2$ decent, and the deformations given by 1-cocycles supported on this cover span $\text{Ann}(V)$.*

Recall from the discussion after Proposition 2.4 that the maps $f_n : X_n \rightarrow X$ are only defined after identifying some points of Σ with each other. The reason we have required all points of Σ to lie in a single element of the open cover is so that we obtain an open cover even after points of Σ are identified with each other.

Proof. The first claim follows because, for n large enough, the map g_n given by Definition 2.2 is close to an isometry for the flat metric.

Let U be the union of the balls of tiny radius centered at points of Σ . We require the radius to be smaller than $\varepsilon/2$. As in the proof of Lemma 9.1, we see that V is supported on $X_n \setminus g_n(X \setminus U)$, which is contained in a single element of the open cover. Hence the (derivatives of) deformations described by cocycles supported on this cover are in $\text{Ann}(V)$.

The dimension of the boundary stratum is the dimension of $\text{Ann}(V)$. The relative homology of (X, Σ) is isomorphic to $H_1(X_n, \Sigma_n)/V$, and the effect of the deformation given by a 1-cocycle on period coordinates is given by this isomorphism. In particular, linearly independent 1-cocycles on X give rise to linearly independent 1-cocycles on X_n .

Thus the cocycles supported on the cover of X_n give rise to a linearly independent set of deformations of dimension equal to the dimension of $\text{Ann}(V)$, and we conclude that these deformations span $\text{Ann}(V)$. \square

Proof of Proposition 9.3. Fix an ε decent cover $\{U_1, \dots, U_p\}$ of the limit surface (X, ω, Σ) such that the derivatives of deformations of (X, ω, Σ) that arise as above from 1-cocycles on this cover span the tangent space to the stratum at (X, ω, Σ) , and so that Σ is contained in a single element of the cover. Pulling back the cover by f_n gives an $\varepsilon/2$ decent cover of $(X_n, \omega_n, \Sigma_n)$. By the previous lemmas, deformations on this cover span in $\text{Ann}(V)$.

There is a neighborhood of 0 in $\text{Ann}(V)$ given by deformations obtained from cocycles of norm at most $\varepsilon/2$. If ξ, ξ_n are in this neighborhood, we have that $(X_n, \omega_n, \Sigma_n) + \xi_n$ and $(X, \omega, \Sigma) + \xi$ are all well defined. It remains to show convergence. We will do this assuming ξ_n, ξ are described by cocycles of norm at most ε' , where $0 < \varepsilon' < \varepsilon/2$ will be specified later.

The complement of the image of g_n is a set of points with small injectivity radius that may have non-trivial topology. We call the complement of the image of g_n the small subsurface of $(X_n, \omega_n, \Sigma_n)$. (There are at most $|\Sigma|$ many connected components of the small subsurface, and there may be fewer.)

$(X_n, \omega_n, \Sigma_n) + \xi_n$ is obtained by gluing the disjoint union of the open sets $f_n^{-1}(U_k)$, with the change in gluing compared to $(X_n, \omega_n, \Sigma_n)$ described by ξ_n . Since the change in gluing does not effect the metric on each element of the cover, and since the small surface is contained in a single element of the cover, it remains in the set of points with small injectivity radius on $(X_n, \omega_n, \Sigma_n) + \xi_n$.

Assume $\varepsilon' < \varepsilon/10$. Then any sufficiently short saddle connection (say, length at most $\varepsilon/100$) on $(X_n, \omega_n, \Sigma_n) + \xi_n$ must be contained in one of these subsurfaces.

To get from $(X, \omega, \Sigma) + \xi$ minus a neighborhood of Σ to $(X_n, \omega_n, \Sigma_n) + \xi_n$ minus the union of the small subsurfaces, it suffices to change the flat metric very slightly on each element of the open cover (g_n is close to an isometry for the flat metric), and then change the gluing instructions very slightly (ξ_n is close to ξ). Hence there should be a map between these two spaces that is almost a flat isometry, and this map verifies the convergence of $(X, \omega, \Sigma) + \xi$ to $(X_n, \omega_n, \Sigma_n) + \xi_n$.

We now sketch one method of constructing such a map. Fix a triangulation of (X, ω, Σ) with vertices Σ . Pick ε' small enough so that this triangulation remains a triangulation on $(X, \omega, \Sigma) + \xi$ whenever $\|\xi\| < \varepsilon'$, that is, all the saddle connections of the triangulation remain saddle connections on $(X, \omega, \Sigma) + \xi$.

Fix $\delta > 0$ be very small, in particular much smaller than the length of the smallest edge in the triangulation. For each edge of the triangulation of (X, ω, Σ) , consider the two points that are distance δ from the endpoints. Fix a given triangle in the triangulation, and denote these points A, B, C, D, E, F as in Figure 9.3 right. These six points are the corners of a hexagon H that shares three edges with the triangle. (One should consider all triangles simultaneously, but for notational simplicity we will focus on just one.)

For n large enough, all such hexagons are contained in the domain of g_n . Furthermore, for n large enough, there is a hexagon H_n spanned by the points $g_n(A), g_n(B), g_n(C), g_n(D), g_n(E), g_n(F)$. The side lengths of H_n converge to those of H .

We may assume that each point on an edge of the triangulation that is distance δ from a vertex is contained in only one element of the open cover, because δ is very small and because of Definition 9.4(1). Each set of the open cover minus a neighborhood of Σ maps continuously

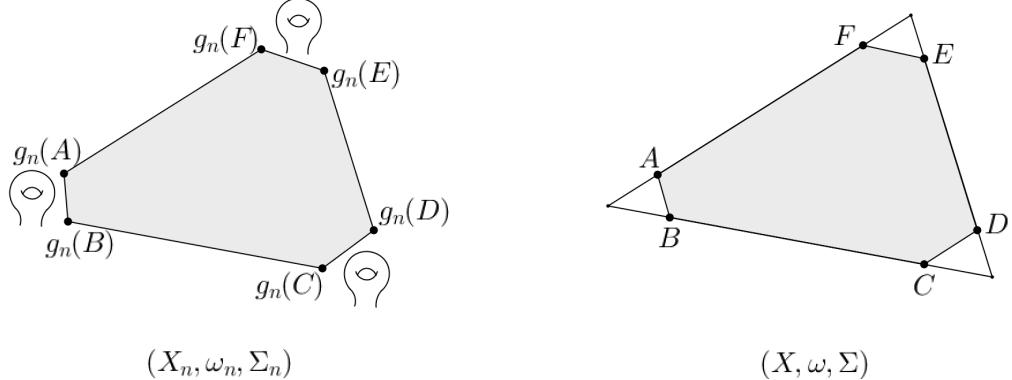


FIGURE 9.3

into both $(X, \omega, \Sigma) + \xi$ and $(X_n, \omega_n, \Sigma_n) + \xi_n$, so we get corresponding points A', B', C', D', E', F' on $(X, \omega, \Sigma) + \xi$ and $A'_n, B'_n, C'_n, D'_n, E'_n, F'_n$ on $(X_n, \omega_n, \Sigma_n) + \xi_n$. For n large enough, these points span hexagons H' and H'_n . As $n \rightarrow \infty$, the side lengths of H'_n converge to those of H' .

Pick a point P in the interior of H' . Joining P to the six points A', B', \dots , we triangulate H' . We can find a piecewise linear map from H' to H'_n that is linear on each triangle, and that has derivative closer and closer to the identity as $n \rightarrow \infty$.

We now have for any fixed δ small enough a sequence of piecewise linear maps from $(X, \omega, \Sigma) + \xi$ minus a neighborhood of Σ depending on δ to $(X_n, \omega_n, \Sigma_n) + \xi_n$ that is closer and closer to preserving the flat metric as $n \rightarrow \infty$. Using the fact that the complement of the image of g_n has injectivity radius going to zero as $n \rightarrow \infty$, we may see that for δ small enough and n large enough, the injectivity radius of the complement of the image of these piecewise linear maps is arbitrarily small. The piecewise linear maps can be approximated by smooth maps that are also close to preserving the flat metric, and a sequence can be chosen along which $\delta \rightarrow 0$ and $n \rightarrow \infty$ to give the sequence required by Definition 2.2 to verify convergence. \square

9.3. Connection to Deligne-Mumford. Here we sketch the proof of the following result. It may be familiar to some experts, and is not used in the paper. The sketch is included for completeness.

Let $\overline{\mathcal{M}}_{g,s}$ be the Deligne-Mumford compactification of the moduli space $\mathcal{M}_{g,s}$ of genus g Riemann surfaces with s marked points. Recall that given a surface (X, ω, Σ) , the points of Σ are the marked points, and Σ is assumed to include the zeros of ω .

Proposition 9.8. $(X_n, \omega_n, \Sigma_n)$ converges to (X, ω, Σ) in the sense of Section 2 if and only if, for each limit point of the sequence in the bundle of Abelian differentials over $\overline{\mathcal{M}}_{g,s}$, the result of removing the components on which the Abelian differential is zero is (X, ω, Σ) .

To be precise, one should remove the components on which the Abelian differential is zero, and the nodes, and then add back in finitely many marked points as required to get a closed surface.

Proof. We will first assume that $(X_n, \omega_n, \Sigma_n)$ converges in the bundle of Abelian differentials, say to (X', ω', Σ') , and argue that this implies convergence as in Section 2. Let U_n be decreasing neighborhoods of the nodes on (X', ω', Σ') , such that $\cap U_n$ is the set of nodes.

We may assume that there is a map g_n from $X' \setminus U_n$ to X_n such that

$$\lim_{n \rightarrow \infty} \omega - g_n^*(\omega_n) = 0.$$

We now claim that the injectivity radius of any point in $X_n \setminus g_n(X' \setminus U_n)$ goes to zero (with respect to the flat metric given by ω_n). This claim follows as in [McM13, Lemma 4.3]: An isoperimetric inequality [McM13, Theorem A.2] gives that, since the annuli in $X_n \setminus g_n(X' \setminus U_n)$ have small boundary, almost all of their area is contained in a cylinder of small circumference. Points in a cylinder of small circumference have small injectivity radius. The boundary of $X_n \setminus g_n(X' \setminus U_n)$ is small because the map g_n is almost isometric.

Remark 9.9. See also [CS09, Theorem 2] for the statement of a more precise isoperimetric inequality, not required here. The work of Rafi [Raf07], [EKZ14, Section 4] provides an alternate approach.

Define G_n as the union of the components of X' on which ω' is nonzero, minus U_n .

We now claim that every point of $X_n \setminus g_n(G_n)$ has injectivity radius going to zero as n goes to infinity. Given the first claim it suffices to show this on the image under g_n of components of X' on which ω' is zero. However, g_n does not distort the flat metric much, so the flat metric is uniformly small on this image. This gives the claim.

Define (X, ω, Σ) to be (X', ω', Σ') with the components where ω' is 0 removed. Convergence $(X_n, \omega_n, \Sigma_n) \rightarrow (X, \omega, \Sigma)$ using the definition in Section 2 follows because the map g_n on (X', ω', Σ') restricted to (X, ω, Σ) gives the map g_n referred to in Section 2.

To show the converse, it suffices to consider $(X_n, \omega_n, \Sigma_n)$ that converge as in Section 2 to (X, ω, Σ) . Without loss of generality, assume they converge to (X', ω', Σ') in the bundle of Abelian differentials over Deligne-Mumford.

Let $(X'', \omega'', \Sigma'')$ be the result of removing the zero area components from (X', ω', Σ') . We need to show that $(X'', \omega'', \Sigma'')$ is equal to (X, ω, Σ) . This is straightforward, because the composition of the map from $(X'', \omega'', \Sigma'')$ to $(X_n, \omega_n, \Sigma_n)$ with the collapse map to (X, ω, Σ) gives maps from a complement of a neighborhood of Σ'' in $(X'', \omega'', \Sigma'')$ to (X, ω, Σ) whose limits in the compact open topology send ω to ω'' . \square

REFERENCES

- [AD16] Artur Avila and Vincent Delecroix, *Weak mixing directions in non-arithmetic Veech surfaces*, J. Amer. Math. Soc. **29** (2016), no. 4, 1167–1208.
- [AG13] Artur Avila and Sébastien Gouëzel, *Small eigenvalues of the Laplacian for algebraic measures in moduli space, and mixing properties of the Teichmüller flow*, Ann. of Math. (2) **178** (2013), no. 2, 385–442.
- [AGY06] Artur Avila, Sébastien Gouëzel, and Jean-Christophe Yoccoz, *Exponential mixing for the Teichmüller flow*, Publ. Math. Inst. Hautes Études Sci. (2006), no. 104, 143–211.
- [AMY13] Artur Avila, Carlos Matheus, and Jean-Christophe Yoccoz, *$SL(2, \mathbb{R})$ -invariant probability measures on the moduli spaces of translation surfaces are regular*, Geom. Funct. Anal. **23** (2013), no. 6, 1705–1729.
- [AN] David Aulicino and Duc-Manh Nguyen, *Rank two affine submanifolds in $\mathcal{H}(2, 2)$ and $\mathcal{H}(3, 1)$* , preprint, arXiv:1501.03303 (2015).
- [ANW16] David Aulicino, Duc-Manh Nguyen, and Alex Wright, *Classification of higher rank orbit closures in $\mathcal{H}^{\text{odd}}(4)$* , J. Eur. Math. Soc. (JEMS) **18** (2016), no. 8, 1855–1872.
- [Api] Paul Apisa, *$GL_2(\mathbb{R})$ orbit closures in hyperelliptic components of strata*, preprint, arXiv:1508.05438 (2015).
- [Bai07] Matt Bainbridge, *Euler characteristics of Teichmüller curves in genus two*, Geom. Topol. **11** (2007), 1887–2073.
- [BCG⁺] Matt Bainbridge, Dawei Chen, Quentin Gendron, Samuel Grushevsky, and Martin Moeller, *Compactification of strata of abelian differentials*, preprint, arXiv:1604.08834 (2016).
- [BCR98] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy, *Real algebraic geometry*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)], vol. 36, Springer-Verlag, Berlin, 1998, Translated from the 1987 French original, Revised by the authors.
- [Boi15] Corentin Boissy, *Connected components of the strata of the moduli space of meromorphic differentials*, Comment. Math. Helv. **90** (2015), no. 2, 255–286.
- [Cal04] Kariane Calta, *Veech surfaces and complete periodicity in genus two*, J. Amer. Math. Soc. **17** (2004), no. 4, 871–908.
- [CE15] Jon Chaika and Alex Eskin, *Every flat surface is Birkhoff and Oseledets generic in almost every direction*, J. Mod. Dyn. **9** (2015), no. 01, 1–23.
- [Che] Dawei Chen, *Degenerations of Abelian Differentials*, preprint, arXiv:1504.01983 (2015).

- [CS09] Jaigyoung Choe and Richard Schoen, *Isoperimetric inequality for flat surfaces*, Proceedings of the 13th International Workshop on Differential Geometry and Related Fields [Vol. 13], Natl. Inst. Math. Sci. (NIMS), Taejön, 2009, pp. 103–109.
- [EKZ14] Alex Eskin, Maxim Kontsevich, and Anton Zorich, *Sum of Lyapunov exponents of the Hodge bundle with respect to the Teichmüller geodesic flow*, Publ. Math. Inst. Hautes Études Sci. **120** (2014), 207–333.
- [EM] Alex Eskin and Maryam Mirzakhani, *Invariant and stationary measures for the $SL(2, \mathbb{R})$ action on moduli space*, preprint, arXiv:1302.3320 (2013).
- [EM01] Alex Eskin and Howard Masur, *Asymptotic formulas on flat surfaces*, Ergodic Theory Dynam. Systems **21** (2001), no. 2, 443–478.
- [EMM15] Alex Eskin, Maryam Mirzakhani, and Amir Mohammadi, *Isolation, equidistribution, and orbit closures for the $SL(2, \mathbb{R})$ action on moduli space*, Ann. of Math. (2) **182** (2015), no. 2, 673–721.
- [EMZ03] Alex Eskin, Howard Masur, and Anton Zorich, *Moduli spaces of abelian differentials: the principal boundary, counting problems, and the Siegel-Veech constants*, Publ. Math. Inst. Hautes Études Sci. (2003), no. 97, 61–179.
- [Fil] Simion Filip, *Semisimplicity and rigidity of the Kontsevich-Zorich cocycle*, preprint, arXiv:1307.7314 (2013).
- [Fil16] ———, *Splitting mixed Hodge structures over affine invariant manifolds*, Ann. of Math. (2) **183** (2016), no. 2, 681–713.
- [Gad] Vaibhav Gadre, *Partial sums of excursions along random geodesics and volume asymptotics for thin parts of moduli spaces of quadratic differentials*, preprint, arXiv:1408.5812 (2014).
- [Gen] Quentin Gendron, *The Deligne-Mumford and the Incidence Variety Compactifications of the Strata of $\Omega\mathcal{M}_g$* , preprint, arXiv:1503.03338 (2015).
- [HLM09] Pascal Hubert, Erwan Lanneau, and Martin Möller, *$GL_2^+(\mathbb{R})$ -orbit closures via topological splittings*, Surveys in differential geometry. Vol. XIV. Geometry of Riemann surfaces and their moduli spaces, Surv. Differ. Geom., vol. 14, Int. Press, Somerville, MA, 2009, pp. 145–169.
- [HLM12] ———, *Completely periodic directions and orbit closures of many pseudo-Anosov Teichmueller discs in $Q(1,1,1,1)$* , Math. Ann. **353** (2012), no. 1, 1–35.
- [HM79] John Hubbard and Howard Masur, *Quadratic differentials and foliations*, Acta Math. **142** (1979), no. 3–4, 221–274.
- [HW] W. Patrick Hooper and Barak Weiss, *Rel leaves of the Arnoux-Yoccoz surfaces*, preprint, arXiv:1508.05363 (2015).
- [IT92] Y. Imayoshi and M. Taniguchi, *An introduction to Teichmüller spaces*, Springer-Verlag, 1992.
- [KM] Abhinav Kumar and Ronen E. Mukamel, *Algebraic models and arithmetic geometry of Teichmüller curves in genus two*, preprint, arXiv:1406.7057 (2014).
- [KZ03] Maxim Kontsevich and Anton Zorich, *Connected components of the moduli spaces of Abelian differentials with prescribed singularities*, Invent. Math. **153** (2003), no. 3, 631–678.

- [LNW] Erwan Lanneau, Duc-Manh Nguyen, and Alex Wright, *Finiteness of Teichmüller curves in non-arithmetic rank 1 orbit closures*, preprint, arXiv:1504.03742 (2015).
- [LW15] Samuel Lelièvre and Barak Weiss, *Translation surfaces with no convex presentation*, Geom. Funct. Anal. **25** (2015), no. 6, 1902–1936.
- [Mas86] Howard Masur, *Closed trajectories for quadratic differentials with an application to billiards*, Duke Math. J. **53** (1986), no. 2, 307–314.
- [McM03] Curtis T. McMullen, *Billiards and Teichmüller curves on Hilbert modular surfaces*, J. Amer. Math. Soc. **16** (2003), no. 4, 857–885 (electronic).
- [McM05] ———, *Teichmüller curves in genus two: discriminant and spin*, Math. Ann. **333** (2005), no. 1, 87–130.
- [McM06a] ———, *Prym varieties and Teichmüller curves*, Duke Math. J. **133** (2006), no. 3, 569–590.
- [McM06b] ———, *Teichmüller curves in genus two: torsion divisors and ratios of sines*, Invent. Math. **165** (2006), no. 3, 651–672.
- [McM07] ———, *Dynamics of $SL_2(\mathbb{R})$ over moduli space in genus two*, Ann. of Math. (2) **165** (2007), no. 2, 397–456.
- [McM13] ———, *Navigating moduli space with complex twists*, J. Eur. Math. Soc. (JEMS) **15** (2013), no. 4, 1223–1243.
- [McM14] ———, *Moduli spaces of isoperiodic forms on Riemann surfaces*, Duke Math. J. **163** (2014), no. 12, 2271–2323.
- [McM15] ———, *Cascades in the dynamics of measured foliations*, Ann. Sci. Éc. Norm. Supér. (4) **48** (2015), 1–39.
- [Möl06a] Martin Möller, *Periodic points on Veech surfaces and the Mordell-Weil group over a Teichmüller curve*, Invent. Math. **165** (2006), no. 3, 633–649.
- [Möl06b] ———, *Variations of Hodge structures of a Teichmüller curve*, J. Amer. Math. Soc. **19** (2006), no. 2, 327–344 (electronic).
- [MT02] Howard Masur and Serge Tabachnikov, *Rational billiards and flat structures*, Handbook of dynamical systems, Vol. 1A, North-Holland, Amsterdam, 2002, pp. 1015–1089.
- [Mum95] David Mumford, *Algebraic geometry. I*, Classics in Mathematics, Springer-Verlag, Berlin, 1995, Complex projective varieties, Reprint of the 1976 edition.
- [MW] Maryam Mirzakhani and Alex Wright, *Full rank affine invariant submanifolds*, preprint, arXiv:1608.02147 (2016).
- [MW02] Yair Minsky and Barak Weiss, *Nondivergence of horocyclic flows on moduli space*, J. Reine Angew. Math. **552** (2002), 131–177.
- [MZ] Martin Möller and Don Zagier, *Modular embeddings of Teichmüller curves*, preprint, arXiv:1503.05690 (2015).
- [Ngu11] Duc-Manh Nguyen, *Parallelogram decompositions and generic surfaces in $H^{\text{hyp}}(4)$* , Geom. Topol. **15** (2011), no. 3, 1707–1747.
- [NW14] Duc-Manh Nguyen and Alex Wright, *Non-Veech surfaces in $H^{\text{hyp}}(4)$ are generic*, Geom. Funct. Anal. **24** (2014), no. 4, 1316–1335.
- [Raf07] Kasra Rafi, *Thick-thin decomposition for quadratic differentials*, Math. Res. Lett. **14** (2007), no. 2, 333–341.
- [SW04] John Smillie and Barak Weiss, *Minimal sets for flows on moduli space*, Israel J. Math. **142** (2004), 249–260.

- [SW07] ———, *Finiteness results for flat surfaces: a survey and problem list*, Partially hyperbolic dynamics, laminations, and Teichmüller flow, Fields Inst. Commun., vol. 51, Amer. Math. Soc., Providence, RI, 2007, pp. 125–137.
- [Vee89] W. A. Veech, *Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards*, Invent. Math. **97** (1989), no. 3, 553–583.
- [Wri14] Alex Wright, *The field of definition of affine invariant submanifolds of the moduli space of abelian differentials*, Geom. Topol. **18** (2014), no. 3, 1323–1341.
- [Wri15a] ———, *Cylinder deformations in orbit closures of translation surfaces*, Geom. Topol. **19** (2015), no. 1, 413–438.
- [Wri15b] ———, *Translation surfaces and their orbit closures: An introduction for a broad audience*, EMS Surv. Math. Sci. **2** (2015), no. 1, 63–108.
- [Wri16] ———, *From rational billiards to dynamics on moduli spaces*, Bull. Amer. Math. Soc. (N.S.) **53** (2016), no. 1, 41–56.
- [Zor06] Anton Zorich, *Flat surfaces*, Frontiers in number theory, physics, and geometry. I, Springer, Berlin, 2006, pp. 437–583.